arXiv:1704.04470v1 [cs.LG] 14 Apr 2017

Lean From Thy Neighbor: Stochastic & Adversarial Bandits in a Network

L. Elisa Celis
elisa.celis@epfl.ch
École Polytechnique Fédérale de Lausanne (EPFL)

Farnood Salehi
farnood.salehi@epfl.ch
École Polytechnique Fédérale de Lausanne (EPFL)

April 17, 2017

Abstract

An individual's decisions are often guided by *those of his or her peers*, i.e., neighbors in a social network. Presumably, being privy to the experiences of others aids in learning and decision making, but how much advantage does an individual gain by observing her neighbors? Such problems make appearances in sociology and economics and, in this paper, we present a novel model to capture such decision-making processes and appeal to the classical multi-armed bandit framework to analyze it. Each individual, in addition to her own actions, can observe the actions and rewards obtained by her neighbors, and can use all of this information in order to minimize her own regret. We provide algorithms for this setting, both for stochastic and adversarial bandits, and show that their regret smoothly interpolates between the regret in the classical bandit setting and that of the full-information setting as a function of the neighbors' exploration. In the stochastic setting the additional information must simply be incorporated into the usual estimation of the rewards, while in the adversarial setting this is attained by constructing a new unbiased estimator for the rewards and appropriately bounding the amount of additional information provided by the neighbors. These algorithms are optimal up to log factors; despite the fact that the agents act independently and selfishly, this implies that it is an approximate Nash equilibria for all agents to use our algorithms. Further, we show via empirical simulations that our algorithms, often significantly, outperform existing algorithms that one could apply to this setting.

1 Introduction

Individuals often have access to information, via their social or economic network, that they can use to make improved decisions. This phenomenon has been observed widely in the social and natural sciences. For instance, a recent work ([Yoo12]) studies farmers who, every year, have to decide which kind of seed to plant (not just what kind of crop, but which variety of seed) in order to attain the most profit (i.e., revenue - cost). In their study, [Yoo12] finds that farmers' decisions are based on a) their own experience in previous years of how different varieties performed, and b) the experiences of peers attained either directly (explicitly via conversations with social contacts) or indirectly (implicitly by observing the farming practices of peers). Moreover, the information farmers used is primarily from peers in their *physical neighborhood* – not only because these are where their contacts are most likely to be, but also because the profit is correlated due to similar soil and weather conditions. These connections between peers then form a network of farmers across the country, where locally, each farmer is trying to learn the best seed for their farm using her own information and that of her neighbors. As another example, consider WI-FI networks in which nodes want to send their data across the best frequency band. Nodes could obtain the current quality of the band their peers are using indirectly through capacity estimation or directly by message passing, and use this information to determine which band to use. Similar social learning phenomena appear in many other areas in various disguises – e.g., in the acquisition

of consumer products by individuals, the adoption of new technologies, the prevalence and spread of corruption, and in the behavior of animals such as squirrels; see [LBG48, KL55, ZAA07, San06, ASC12]. A natural question then is: *How should an individual incorporate the information from their neighbors in order to make the best decisions, and how much improvement can such information bring?*

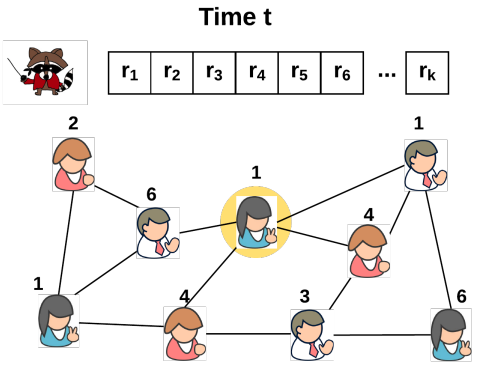

Figure 1: A social network in which all individuals play against the same bandit, i.e., if two individuals select the same arm at the same time step, they observe the same reward (up to noise). At each time t, each individual selects an arm (shown), and then observes its reward along with the actions/rewards of her neighbors. E.g., the yellow circled individual would observe the reward of actions 1, 4 and 6 in this time step.

Consider the following formulation geared towards capturing the type of settings mentioned above (see also Figure 1): at each time step each individual selects one of K possible actions, observes the value or reward of selecting that action, and observes the actions, values and/or decision process of their neighbors in the network. This selection and observation is repeated again and again, and each individual has the end goal of identifying the action $a^\star \in K$ that brings them the best value over all time steps; i.e., minimizing the *regret*. This formulation seems to suggest that the problem is suited for study using the bandit optimization framework, except that now there is *additional information available to an individual via her neighbors*.

Towards this, one approach could be to consider the framework of bandits with *side observations* for which, in the adversarial setting, variants of the multiplicative-weight update algorithm have been developed with success. Informally, side observations just mean that at each time step, in addition to observing the reward of a selected action $a(t)$, one may observe (but not receive) rewards from a set of other actions $S(t)$. A recent body of work has explored how to minimize regret for various different models of $S(t)$. In the *free observation model*, the individual is allowed to select $S(t)$ up to some cardinality (e.g., [AKTT15]). However, if one tried to apply such algorithms to the social settings considered above, it would require an individual to decide which actions her neighbors should take, and hence is not feasible as a solution in this setting. In another line of work (e.g., [ACBDK15]) an *action-network model* has been studied: Here, the actions form a network and the individual observes the rewards of the neighbors of the *action* she selects (as opposed to the rewards of the actions that *her* neighbors select). The action-network is often taken to be exogenous and can be changing over time. Thus, one can apply the algorithms developed in the action-network setting to the social setting above by defining $S(t)$ to be the set of actions selected by the individual's neighbors; however, this may not always be optimal for the social setting as neighbors can provide even more information (see Section 4.5).

The above approaches have been developed in independent contexts and hence geared towards different settings. Towards obtaining optimal results in the social setting described above, the challenge is to adequately model the information from neighbors that can aid in learning, leverage it appropriately, and quantify the advantage it provides. In this paper we present such a model and show how one can incorporate the additional social information in order to obtain optimal results. More specifically, in the stochastic setting we show that incorporating side information in a simple way gives rise to a near-optimal algorithm. Furthermore, in the adversarial setting, we present a modified multiplicative-weight update algorithm that uses an new unbiased estimator to incorporate this information appropriately into the estimation of the reward of each action. The proofs requires us to overcome some additional hurdles in order to bound the amount of information gleaned from the neighbors and attain optimal bounds on the regret. Our algorithms outperform other state-of-the-art bandit algorithms both theoretically (Section 4) and empirically (Section 5).

1.1 Summary of Our Results

Since our model appeals to the bandit framework, we start by recalling its salient features: In the bandit optimization setting, there are K potential actions (*arms*), and the individual selects one arm at each time step. Each arm j has an

(unknown) *reward* $g_j(t)$ at *time* $t = 1, \ldots, T$. The individual receives the reward $g_{a(t)}(t)$ for the selected arm $a(t)$, while the rewards $g_j(t)$ for all other actions $j \neq a(t)$ remain unknown. Ideally, the individual would like to select the arm with the best overall reward, i.e., $j^\star = \arg\max_j \mathbb{E}\left[\sum_{t=1}^T g_j(t)\right]$. However, in the absence of knowledge about the reward structure, this is not feasible. Instead, as is prevalent in the online learning literature, $j^\star$ is used as a *benchmark* and the *regret* R (i.e., the difference between the individual's rewards and those of the best arm) is minimized. Formally,

$$R \stackrel{\text{def}}{=} \mathbb{E}\left[\sum_{t=1}^T g_{j^\star}(t) - \sum_{t=1}^T g_{a(t)}(t)\right],$$

where the expectation is taken over the randomness in the rewards (if stochastic) and the randomness in the algorithm.[1] An important divide arises from how the rewards are decided: for *stochastic bandits*, the rewards $g_j(t)$ are i.i.d. from an (unknown) distribution $\mathcal{F}_j$, and for *adversarial bandits* (also known as non-stochastic bandits), the rewards $g_j(t)$ are set in an arbitrary manner by an *adversary* that knows the individual's algorithm and past coin flips. In realistic settings, such as the examples mentioned in the introduction, the situation is often somewhere in-between. Thus, in this work we consider both extremes by analyzing both the stochastic and adversarial settings. In either case, algorithms must carefully tradeoff between *exploration* (gaining new knowledge about the rewards) and *exploitation* (using current knowledge to maximize rewards) in order to minimize regret. In either case, algorithms must carefully tradeoff between *exploration* (gaining new knowledge about the rewards) and *exploitation* (using current knowledge to maximize rewards) in order to minimize regret.

The Model. Firstly, we assume all individuals play against the same multi-armed bandit: in the stochastic setting, the reward distributions $\mathcal{F}_i$ is the same for arm i for all individuals (although the realizations at any given time may differ), and in the adversarial setting the reward vector $\mathbf{g}(t)$ selected by the adversary at time t is the same for all individuals. Clearly, there must be some similarity in the rewards between neighbors for social learning to occur. Our results also extend to the setting in which the distributions or rewards are correlated, e.g., via 0-mean noise (i.e., each individual i receives the reward as above + individual noise); for ease of presentation we omit this extension, the proofs follow analogously. Secondly, we assume that each individual can observe the following for all *neighbors* i:

(1) the actions $a_i(t)$,
(2) the rewards $g_{a_i(t)}(t)$, and
(3) (for the adversarial setting only) the probability distribution each neighbor used to select an arm at *the previous* time step.

Assumptions (1) and (2) are natural and directly inspired by applications such as those mentioned in the introduction; individuals either directly or indirectly observe their neighbors' actions and rewards. (3) additionally assumes a limited knowledge of how neighbors made their decisions on a step-by-step basis, without having to assume we know their overall algorithm or restricting their behavior in any way. All individuals are free to select their probability distributions arbitrarily (and depend arbitrarily on each other), and each one can draw her decision independently of the rest. While it would be nice to drop assumption (3) entirely, this would prevent us from attaining optimal regret bounds (see Proposition 4.5 and the discussion in Appendix A.4).[2]

Importantly, an individual

(4) *does not* know about the actions and rewards of individuals beyond her neighbors,
(5) *does not* know any global properties of the network,
(6) *does not* know which algorithm other individuals (including neighbors) are using, and
(7) *cannot* dictate or coerce other individuals to act a certain way.

Removing any of assumptions (4)-(7) would be unnatural for the social learning setting described above: If (4) does not hold, we would simply consider such an individual a neighbor, removing (5) is unnatural as the network can be very large and we cannot expect to have knowledge of distant individuals, removing (6) seems impractical as it would mean that individuals have a detailed knowledge of how neighbors select actions, and allowing individuals to be coerced as in

[1] This is often referred to as *pseudo-regret*; in this paper we simply refer to it as *regret*.

[2] Alternatively, under different assumptions (e.g., if we assume the neighbors are using an algorithm such as EXP3) we can estimate these distributions which would suffice.

(7) would be in conflict with the idea that every individual seeks to improve her own performance. Hence, given (1)-(7), any improvement in the individual's regret arises solely from *passive* observation of *local* information.

Remark 1.1. *We emphasize that, by setting up the model as above, agents can act independently and work to selfishly minimize their own regret. Hence, (as our algorithms are near-optimal) it is always an individual's (approximate) best-response to use our algorithms,* regardless of what her neighbors do*. Putting it another way, when the action space is the set of algorithms, it is an (approximate) Nash equilibrium for all nodes to use ours – no individual can significantly improve their regret by deviating to use an alternate algorithm. This gives rise to an interesting set of questions regarding the average regret of a network (depending on its structure) in equilibrium, or the expected regret in equilibrium of a node as a function of its position in the network. Some immediate implications about the properties of such equilibria follow directly from our results.*

The Algorithm in the Stochastic Setting. Algorithms for the classical stochastic bandit setting [AG12, ACBF02, GC11] maintain metrics about the observed samples, and determine which arm to select based on these metrics. Hence, a natural strategy would be for an individual in the networked setting to take one of these algorithms and incorporate samples obtained from neighbors along with her own without differentiating between the two. Indeed, we this simple insight suffices to get near-optimal results. Our UCBN algorithm extends the classic UCB algorithm by incorporating all observed samples indiscriminately. We show that this suffices to improve performance, often dramatically, both asymptotically and in silico. These results are presented in Section 3 and Appendix B. We show that the regret of our algorithm interpolates between $O(1)$ and the $O(K \ln T)$ regret for the classic bandit setting depending on the amount of exploration conducted by the neighbors (see Theorem 3.1), and these bounds are asymptotically optimal (see Theorem B.3). As a corollary, in the complete network with b vertices, if all neighbors use our algorithm (i.e., in equilibrium) the regret reduces to $O(\frac{K \ln T}{b})$, which is optimal as even a fully centralized approach can speed up learning by at most a factor of the increase in the number of available samples, i.e., b (see Corollary B.2). The theoretical results are presented in Section 3 and Appendix B), and the empirical results in Section 5.2).

The Algorithm in the Adversarial Setting. Algorithms for classic adversarial bandits are, typically, variations of the multiplicative weights update method. Such algorithms maintain a weight for each arm, and (multiplicatively) update that weight based on the observed reward(s). An individual then selects an action proportionally to the weight vector. It is easy to verify that if we naïvely incorporate samples from neighbors as if they were our own into the weight vector there is no improvement in the regret. Our variation of the multiplicative-weights update algorithm a) incorporates information in a clever manner (that allows for new improved bounds) by using a different unbiased estimator for the rewards and then b) adapts according to the behavior of its neighbors by tuning its update parameters. We show that the regret of our algorithm provably interpolates between the $O(\sqrt{KT \ln K})$ regret for the classical bandit setting and the $O(\sqrt{T \ln K})$ regret for the full-information setting (where the full vector $\mathbf{g}(t)$ is observed at each time step) depending on the amount of exploration conducted by the neighbors (see Theorem 4.1), and is optimal up to log factors (see Theorem 4.2 & 4.5). Moreover, our algorithm improves performance over the state-of-the art in silico. The theoretical results are presented in Section 4 and Appendix A), and the empirical results in Section 5.1).

2 Related Work

Distributed learning in a network is a broad topic and has been studied under various names in several disciplines. However, to the best of our knowledge, our model for learning from neighbors, along with its assumptions and non-assumptions (1)-(7) which are motivated by relevant settings in sociology and economics, is novel. Here we briefly survey the closest relatives to our work.

In the study of *non-strategic learning on networks*, individuals are connected via a network, and each individual has a finite set of actions with probabilistic rewards whose distributions depend on the *state of the world* (see [Goy05], Chapter 2 for a survey). Indeed, would be similar to our model in the stochastic setting. However, work in this area has focused on studying variants of a *greedy* algorithm, and answering the question of whether learning (i.e., discovery of the state of the world, and hence convergence to the best action) occurs asymptotically (see, e.g., [BG98, EF93, BG01, GK03, GJ10]). Instead, we are concerned with *regret*, which could be loosely interpreted as the *rate of convergence*.

Recall that in models of bandits with *side observations*, in addition to observing $g_{a(t)}(t)$, one may *observe* (but not receive) additional rewards $g_{S(t)}(t)$. The set of arms $S(t)$ depends on the particular model of side observations. In the free observations model, the individual can select B additional arms to observe at each time step; i.e., $|S(t)| = B$ and the individual selects $S(t)$ for all t. Such models have been studied both for stochastic ([YM09]) and adversarial ([AMS12, AKTT15]) bandits. Without assumption (7), we could apply such algorithms directly because an individual could dictate which actions her neighbors should take. In the social setting we cannot hope to control our neighbor's decisions in this manner. Still, we show that the performance of our algorithm is equivalent empirically to such algorithms (see Section 5). In the arm-network (or action-network) setting, the individual observes the rewards of the neighbors of the arm she selects. Such stochastic ([CKLB12, BES14]) and adversarial ([MS11, ACBGM13, KNVM14, ACBDK15]) bandit settings have been studied. While one could apply these algorithms to the social setting, some social information, in particular from assumption (3), is left on the table. Leveraging this allows us to provably outperform such approaches (see Section 4.5), and empirically the difference can be dramatic (see Figure 3).

Other work has considered bounding the *cumulative* regret of all individuals, rather than individuals minimizing their own regret. Towards this, centralized algorithms for various versions of stochastic bandits have been studied, in particular for the complete graph ([BES13, SBFH$^+$13, CBGZ13]). Although the centralized setting is not the object of our study, as a corollary, we obtain a centralized algorithm for adversarial bandits that is optimal on the complete network (see Section 4.6).

3 Technical Overview for the Stochastic Setting

To describe our algorithm, let us first revisit the UCB algorithm, first introduced by Auer et. al. [ACBF02] and since widely extended and studied (see, e.g., [Bub10, MMS11, GC11]). UCB is an asymptotically optimal algorithm for stochastic bandits with many well-studied variants (see, e.g., [BCB12] for an overview). The idea behind the algorithm uses the principle of *optimism in the face of uncertainty*; the algorithm maintains an optimistic *upper bound* on the mean reward of each arm, and selects an arm with maximal upper bound. As is standard, we assume the probability distributions satisfy Hoeffding's lemma. Then, arm j at time t has an upper bound

$$U_j(t) = \widehat{\mu}_j(t) + \sqrt{\frac{\alpha \ln(t)}{2n_j(t)}} \tag{1}$$

which holds with probability at least $1 - t^{-\alpha}$ when $\widehat{\mu}_j(t)$ is the sample mean of arm j, and $n_j(t)$ is the number of samples we have for j. At time t, an arm $a(t) \in \arg\max_j\{U_j(t)\}$ is selected.

We make a simple extension to UCB for an agent on a network: the agent simply incorporates *all* samples and *all* rewards into n_j and $\widehat{\mu}_j$ regardless of whether it came from her action or was observed from one of her neighbors. We denote this algorithm by UCBN, and note that it can be implemented by an individual irrespective of the graph structure and the algorithm(s) her neighbors may employ. The regret of UCBN algorithm is upper bounded as follows.

Theorem 3.1. *Consider an agent with neighbors who play arbitrarily. Let $n'_i(t)$ be the number times arm i has been selected by one of her neighbors by time t. Then, the regret of* UCBN *for any $\alpha > 2$ is*

$$R \leq \sum_{i,\Delta_i>0} \left(\max\left\{ \max_{t=1,..,T} \left\{ \frac{2\alpha \ln t}{\Delta_i} - n'_i(t)\Delta_i \right\}, 0 \right\} + \frac{\alpha}{\alpha - 2} \right), \tag{2}$$

where Δ_i is the difference between $\mu_{i^\star}$ and μ_i.

This result is asymptotically optimal (see Theorem B.3). The regret differs from the regret of the classic UCB regret by the $-n'_j(T)\Delta_j$ term, and, depending on the behavior of the neighbors, can potentially take the agent from logarithmic to constant regret.

Clearly, the performance of an agent must depend on the behavior of her neighbors. In the worst case, if there are *clumsy* agents who always select the same arm, then our regret is not improved much. However, as long as the agent has at least one neighbor who explores an arm uniformly at random with probability $\varepsilon_t \in \omega(\frac{K \ln t}{t})$ at time t (e.g., this occurs if a neighbor uses an adaptive greedy algorithm), then the regret is $O(1)$! Hence, this allows us to

interpret neighbor behavior to our regret seamlessly. As an instructive example, consider the setting where all agents use UCBN in a complete graph. The regret in this setting is $O\left(\frac{K \ln T}{b}\right)$. In other words, the regret of an agent using UCBN is a factor $O(1/b)$ less than that of an agent using UCB – indeed we cannot hope to do better, even in a completely centralized setting. The proof parallels the proofs for the original UCB results (see, e.g., [BCB12] for a template), and can be found along with further discussion in Appendix B. While the story for the stochastic setting turns out to be simple and easy to manage, the adversarial setting, as we see below, turns out to be more challenging.

4 Technical Overview for the Adversarial Setting

4.1 Preliminaries

The multiplicative weight update method has been discovered many times in many fields over the past century (see [AHK12] for an overview). It is a simple yet surprisingly powerful way to *conservatively* update beliefs about the benefit of a given arm and is extremely effective for adversarial bandits and is asymptotically optimal up to log factors (see, e.g., [ACBFS02, FKM05, ACBGM13]). Such algorithms for the full information setting (where all rewards are observed at each time step) maintain a vector of weights w_j for each arm j, and (multiplicatively) update it at each time step: $w_j(t+1) = w_j(t)e^{\delta g_j(t)}$, where $0 \leq g_j(t) \leq 1$ is the *reward* observed for arm j at time t and δ is the *update parameter*. The probability of choosing arm j at time t is proportional to the weight $w_j(t)$, namely, $p_j(t) = \frac{w_j(t)}{W_t}$ where $w_j(0) = 1/N$, and $W_t = \sum_{j=1}^{N} w_j(t)$. This algorithm, for an optimal choice of δ has regret $\Theta(\sqrt{T \ln K})$. Extending to the bandit setting uses a simple trick; instead of using $g_j(t)$ to update, we use an *unbiased estimator* $\widehat{g}_j(t)$ for $g_j(t)$ (see, e.g., [ACBFS02, FKM05]). Since we no longer observe the reward of all of the arms, we must also ensure some exploration should be added to the algorithm. This is achieved by setting a lower bound $\eta \in [0, 1]$ (the exploration parameter) on the probability of exploration: $p_j(t) = (1-\eta)\frac{w_j(t)}{W_t} + \frac{\eta}{K}$. This algorithm, also known as EXP3 [ACBFS02], achieves regret $O(\sqrt{TK \ln K})$, and is optimal up to log factors for the right choice of parameters (see [BCB12] for an exposition).

4.2 Formal Statement of Results

We call our algorithm in the adversarial setting EXPN. Recall that $p_j(t)$ is the probability that an individual selects arm j at time t. Let $q_j^i(t)$ be the probability that her neighbor i selects arm j at time t. We denote the number of an individual's neighbors by b. Note that the number of nodes in a network, denoted by N, may be much larger, but the remaining network does not play a role in the algorithm or main results.

Theorem 4.1. *Given an individual with b neighbors who are playing arbitrarily, the regret when using the* EXPN *algorithm is*

$$R_{\mathrm{EXPN}} \leq \mathbb{E}\left[2\sqrt{\left(T + \sum_{t=1}^{T} \gamma_t\right) \ln K}\right] \tag{3}$$

where $\gamma_t = \sum_{j=1}^{K} \frac{p_j(t)}{p_j(t) + \sum_{\ell=1}^{b} q_j^\ell(t)}$.

Before discussing the proof, we first re-write the results in a way that makes them easier to interpret.

For ease of presentation, momentarily assume that for all arms $j \in [K]$, neighbors $i \in [b]$, and times $t \in [T]$ we have that $q_j^i(t) \geq \frac{\varepsilon_i}{K}$ for some $\varepsilon_i \in (0, 1]$.[3] We can then reinterpret R_{EXPN} as a function of the bandit regret (R_{EXP3}) and full information information regret (R_{FULL}) as follows:

$$R_{\mathrm{EXPN}} = \begin{cases} R_{\mathrm{FULL}} \cdot \sqrt{\beta} & \Theta \leq 1 - \sqrt{\frac{\ln K}{\beta T}} \\ R_{\mathrm{EXP3}} \cdot \sqrt{\frac{\beta + K/\Theta}{2K}} & 1 - \sqrt{\frac{\ln K}{(\beta + \frac{K}{\Theta})T}} \leq \Theta \end{cases} \tag{4}$$

[3] This assumption is not required for the proof of Theorem 4.1, only for the ease of interpretation in Equation 4. Note that if a neighbor is running any variant of the multiplicative weight update method, this condition is satisfied. Removing this assumption requires the *number* of non-zero ε_j to be tracked for each j, and these numbers would appear in the regret bound.

where $\Theta = \Pi_{i=1}^{b}\left(1-\frac{\varepsilon_i}{K}\right)$ and $\beta = \frac{1}{1-\left(1-\frac{1}{K}\right)\Theta}+1$.
In particular, note that when $\Theta = 1$, none of the individual's neighbors maintain a probability distribution that is bounded away from 0 for all arms. In other words, the neighbors are not exploring effectively. In this case, $\beta = K+1$ and $R_{\text{EXPN}} \in O(\sqrt{TK\ln K})$, the same as in the classical bandit setting. On the other hand, for example, when $\Theta \leq 1/2$, then $\beta \leq 3$ and hence $R_{\text{EXPN}} \in O(\sqrt{T\ln K})$, the same as in the full-information setting. Hence, this algorithm smoothly interpolates between bandit regret and full information regret as a function of the neighbors' exploration.

The proofs of Theorem 4.1 and Equation 4 appear in Section 4.4 and Appendix A respectively. At first, the proofs parallel standard approaches to analyze the multiplicative-weight update method; the crucial difference is a new unbiased estimator that is used in order to incorporate the neighbors' information (see Section 4.3). This leads to the following bound on the regret:

$$R \leq \frac{\ln K}{\delta} + \eta T + \delta T \sum_{j=1}^{K} \frac{p_j(t)}{p'_j(t)}.$$

The technical obstacle then becomes attaining tight bounds on the $\sum_{j=1}^{K}\frac{p_j(t)}{p'_j(t)}$ term (see Lemmas A.2 and A.4). What remains is then a straightforward (albeit tedious) optimization problem on the parameters η and δ.

We can further show that Theorem 4.1 is optimal up to log factors.

Theorem 4.2. *Given an individual with b neighbors who are playing arbitrarily, the regret when using the* EXPN *algorithm is*

$$R_{\text{EXPN}} = \Omega\left(\sqrt{T + \sum_{t=1}^{T}\gamma_t}\right)$$

where $\gamma_t = \sum_{j=1}^{K}\frac{p_j(t)}{p_j(t)+\sum_{\ell=1}^{b}q_j^{\ell}(t)}$ as defined above.

This information theoretic lower bound follows by extending the lower bound attained for the classic bandit setting (see, e.g., [ACBF02]), and is given in Appendix A.3. This theorem shows that the analysis of our algorithm is tight up to log factors and will help us establish dominance over other potential approaches as discussed in Section 4.5 – a weaker lower bound that is not algorithm-dependent (which matches the above bound for pathological cases such as when all neighbors always play the same action) is presented in Theorem 4.5.

4.3 The EXPN Algorithm

Key to our EXPN algorithm is the following new unbiased estimator for the rewards:

$$\widehat{g}_j(t) = \begin{cases} \frac{g_j(t)}{p'_j(t)} & \text{if some individuals selects action j at time t} \\ 0 & \text{otherwise,} \end{cases} \tag{5}$$

where $p'_j(t)$ is the probability that at least one individual selects action j, i.e.,

$$p'_j(t) \stackrel{\text{def}}{=} 1-(1-p_j(t))(1-q_j^1(t))\cdots(1-q_j^b(t)). \tag{6}$$

The algorithm then updates the weights according to $w_j(t+1) = w_j(t)e^{\delta\widehat{g}_j(t)} = w_j(0)e^{\delta\sum_{s=1}^{t}\widehat{g}_j(s)}$, where $w_j(0)=1$, and updates the probability distributions according to $p_j(t) = (1-\eta)\frac{w_j(t)}{W_t}+\frac{\eta}{K}$ where $W_t = \sum_j w_j(t)$. Note that this algorithm can be implemented irrespective of the network structure and depends only on the information obtained locally from neighbors as defined in our model. In essence, the key to our algorithm is two fold:

1. *Design* a new unbiased estimator $\widehat{g}_j(t)$ that incorporates the side observations obtained from neighbors: Unlike for stochastic bandits, naïve estimators do not suffice, and a new approach is required.[4]

[4]We must ensure that in bounding $\mathbb{E}[\widehat{g}_j^2]$, we get some improvement over the usual bandit setting; it is easy to verify that such bounds do not hold for naïve estimators such as the average of the neighbors' estimators.

2. *Decouple* the exploration and exploitation parameters: When an individual's neighbors explore a lot, she could benefit by *free-riding* off of the exploration of her neighbors; this is accomplished by decreasing her exploration parameter η. However, if we take $\delta = \eta/K$ as in EXP3, this dampens our updates. Hence we need η and δ to act independently.

This analysis is presented in Appendix A.1.

4.4 Proof of Theorem 4.1

While the above version of the algorithm gives a natural interpretation of the parameters, in order to attain the stronger regret bound in Theorem 4.1 we take a slightly different approach. Instead of decoupling the parameters η and δ, we instead use an adaptive δ_t that changes based on the amount of information received from the neighbors.[5] In particular, we let $\delta_t = \sqrt{\frac{\ln K}{\sum_{c=1}^{t}(1+\gamma_c)}}$ and can take $\eta = 0$. Importantly, note that we use the same unbiased estimator in either version of the algorithm.

Proof of Theorem 4.1. The first part of proof (from Equation (7) to Equation (12)) parallels the traditional analysis for analyzing multiplicative weight update algorithms; for completeness we present the steps without going into the details (see [BCB12] for an exposition). We first write $\mathbb{E}[\widehat{g}_j(t)]$ as follows:

$$\mathbb{E}[\widehat{g}_j(t)] = \frac{1}{\delta_t}\left(\ln \mathbb{E}\left[\exp\left(-\delta_t\left(\widehat{g}_j(t) - \mathbb{E}[\widehat{g}_j(t)]\right)\right)\right] - \ln \mathbb{E}\left[\exp\left(-\delta_t \widehat{g}_j(t)\right)\right]\right) \tag{7}$$

where the expectation is over the randomness of the estimator and choice of the arm. We will now consider the right-hand side of the equation and upper bound the two terms separately.

$$\begin{aligned}\frac{1}{\delta_t}\ln \mathbb{E}\left[\exp\left(-\delta_t(\widehat{g}_j(t) - \mathbb{E}[\widehat{g}_j(t)])\right)\right] &= \frac{1}{\delta_t}\ln \mathbb{E}\left[\exp(-\delta_t\widehat{g}_j(t)] + \mathbb{E}[\widehat{g}_j(t)]\right) \\ &\leq \frac{1}{\delta_t}\mathbb{E}[\exp(-\delta_t\widehat{g}_j(t)) - 1 + \delta_t\widehat{g}_j(t)] \leq \frac{\delta_t}{2}\mathbb{E}[\widehat{g}_j^2(t)],\end{aligned} \tag{8}$$

where we use the inequalities $\ln x \leq x - 1$ and $\exp(-x) - 1 + x \leq x^2/2$ for $x \geq 0$. Now, let $\widehat{G}_j(t) = \sum_{a=1}^{t}\widehat{g}_j(t)$, and let $\psi(t) = \frac{1}{\delta_t}\ln\left(\frac{1}{K}\sum_{j=1}^{K}\exp(-\delta_t\widehat{G}_j(t))\right)$. Hence,

$$\begin{aligned}-\frac{1}{\delta_t}\ln \mathop{\mathbb{E}}_{a_t\sim p'_t}\mathop{\mathbb{E}}_{j\sim p'_t}\exp(-\delta_t\widehat{g}_j(t)) \leq -\frac{1}{\delta_t}\mathop{\mathbb{E}}_{a_t\sim p'_t}\ln \mathop{\mathbb{E}}_{j\sim p'_t}\exp(-\delta_t\widehat{g}_j(t)) \\ -\frac{1}{\delta_t}\mathop{\mathbb{E}}_{a_t\sim p'_t}\left[\ln\left(\sum_{j=1}^{K}\frac{\exp(-\delta_t\widehat{G}_j(t))}{\sum_{c=1}^{K}\exp(-\delta_t\widehat{G}_c(t-1))}\right)\right] \leq \psi(t-1) - \psi(t)\end{aligned} \tag{9}$$

where we recall that $\eta = 0$, the first inequality we use Jensen's inequality, and note that δ_t is decreasing in t. By summing up the terms in Equation (7) and (9) over all t we get

$$\sum_{t=1}^{T}\mathbb{E}[\widehat{g}_j(t)] \leq \sum_{t=1}^{T}\frac{\delta_t}{2}\mathbb{E}[\widehat{g}_j^2(t)] - \mathop{\mathbb{E}}_{a_t\sim p'_t}\psi(T). \tag{10}$$

We now bound $-\psi(T)$ as follows:

$$-\psi(T) = \frac{\ln K}{\delta_T} - \frac{1}{\delta_T}\ln\left(\sum_{j=1}^{K}\exp(-\delta_T\widehat{G}_j(T))\right) \leq \frac{\ln K}{\delta_T} - \frac{1}{\delta_T}\ln\left(\exp(-\delta_T\widehat{G}_k(T))\right) = \frac{\ln K}{\delta_T} + \widehat{G}_k(T). \tag{11}$$

[5] Adaptive δ_t are often used when T is unknown; here the adaptive δ_t serves a different function by allowing us to explicitly respond to the neighbors' actions.

Plugging this into in the above inequality yields

$$\sum_{t=1}^{T} \mathbb{E}[\widehat{g}_j(t)] \leq \sum_{t=1}^{T} \frac{\delta_t}{2} \mathbb{E}[\widehat{g}_j^2(t)] + \frac{\ln K}{\delta_T} + \mathop{\mathbb{E}}_{a_t \sim p'_t}[\widehat{G}_k(T)]. \tag{12}$$

Now, we note that given $p_j(t)$ at time t we have

$$\mathbb{E}\left[\widehat{g}_j(t)\right] = g_j(t), \tag{13a}$$

$$\mathbb{E}\left[\sum_{t=1}^{T}\sum_{i=1}^{K} p_j(t)\widehat{g}_j(t)\right] = \sum_{t=1}^{T}\sum_{i=1}^{K} p_j(t)g_j(t) = \mathbb{E}\left[\sum_{t=1}^{T} g_{a(t)}(t)\right], \text{ and} \tag{13b}$$

$$\mathbb{E}\left[\sum_{t=1}^{T}\sum_{i=1}^{K} p_j(t)\widehat{g}_j^2(t)\right] = \sum_{t=1}^{T}\sum_{i=1}^{K} \frac{p_j(t)}{p'_j(t)} g_j^2(t) \leq \sum_{t=1}^{T}\sum_{i=1}^{K} \frac{p_j(t)}{p'_j(t)} \tag{13c}$$

as $g_j(t) \leq 1$, and where the expectation is over randomness of the algorithm. Now, attaining a good bound on the regret boils down to attaining a good bound on $\sum_{i=1}^{K} \frac{p_j(t)}{p'_j(t)}$. Towards this, we need a technical lemma that, in effect, allows bounds the amount of information received from neighbors.

Lemma 4.3. $\sum_{j=1}^{K} \frac{p_j}{1-(1-p_j)(1-q_j^1)\cdots(1-q_j^b)} \leq \sum_{j=1}^{K} \frac{p_j}{p_j+q_j^1+q_j^2+\cdots+q_j^b} + 1.$

The proof is presented in Appendix 4.4. Using this Lemma and combining all of the above, we get

$$R \leq \frac{\ln K}{\delta_T} + \frac{1}{2}\sum_{t=1}^{T} \delta_t(1+\gamma_t). \tag{14}$$

To conclude the proof, recall that $\delta_t = \sqrt{\frac{\ln K}{\sum_{c=1}^{t}(1+\gamma_c)}}$, use Lemma 3.5 of [ACBG02], and take expectation of the both sides of Equation (14). □

4.5 Comparison to Alternate Approaches

Instead of developing a new algorithm, we could have attempted to leverage an existing one. The most natural one to try is from the arm-network setting which is as follows: there is a single individual and the bandit's arms form an arm-network which can change over time. An edge from arm u to arm v means that by choosing arm u we observe the reward of arm v. Thus we could, in retrospect at each time step, recreate an arm-network (see Figure 2) and apply an arm-network algorithm. We consider EXP3G ([ACBDK15]), which is the state-of-the-art solution for such problems, and performed best amongst arm-network algorithms in our empirical simulations. However, we prove in Appendix A.4 that our algorithm is at least as good.

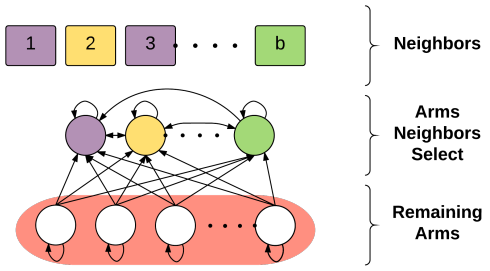

Figure 2: The arm-network has an edge from arm u to arm v if, having selected arm u, we observe the reward of arm v. Arms selected by neighbors in the social network are form a clique, and the remaining arms have self loops and edges to all selected arms.

Proposition 4.4. $R_{\text{EXPN}} = O(R_{\text{EXP3G}})$.

Moreover, as we will see in Section 5, the regret of EXPN is often drastically better empirically. This because EXP3G, and other similar algorithms, were developed from different settings in which it was not possible to make use of the probability distributions afforded to us by assumption (3). Indeed, without this assumption, one can get a stronger lower bound than the one presented above.

Proposition 4.5. *Let n_t be the size of the set of arms selected (arbitrarily) by all of the individual's neighbors at time t. Then, the regret $R_{\mathcal{A}}$ for any algorithm $\mathcal{A}$ in our setting without assumption (3) is* $R_{\mathcal{A}} = \Omega\left(\sqrt{T + \sum_{t=1}^{T}(K - n_t)}\right)$.

The proof follows from Theorem 5 of [ACBG$^+$14]. Our EXPN algorithm is often able to beat this bound by leveraging (3). For example, this proposition implies that if we have a complete network on b vertices where $\log K \ll b \ll K$, then $R_{\text{EXP3G}} = \Omega\left(\sqrt{(K-b)T}\right)$ while in our case $R_{\text{EXPN}} = O\left(\sqrt{\frac{K}{b}T}\right)$ (see Corollary 4.7).

4.6 A Centralized Solution for the Network

Our model and algorithm are formulated for an individual because this allows us to draw the most general conclusions – bounding the individual's regret as a *function* of the neighbors' behavior. However, a surprising feature is that it can also be made into a centralized solution. In the general case, this requires assuming there is an external coordinator that can select a maximum-degree individual to lead and direct the rest on how to act as follows: Let $v^\star$ be the maximum degree node selected. The coordinator directs $v^\star$ to use the EXPN algorithm. The remaining nodes u are each assigned a neighbor v_u that lies on the shortest path between them and $v^\star$, and are directed to copy the probability distribution that v_u used in the previous time step.

Theorem 4.6. *Using the above centralized algorithm, the regret of* all *individuals is at most*

$$R = O\left(\Delta + \sqrt{\left(1 + \frac{K}{1+b_{max}}\right) T \ln K}\right), \tag{15}$$

where b_{max} is the degree of $v^\star$ and Δ is the diameter of the network.

The proof follows, with minor modifications, from the proof of Theorem 4.1; the main difference regards accounting for the delay (of at most Δ time steps) for the farthest node from $v^\star$ to update their probability distribution. By replacing γ_t with $\frac{K}{1+b_{max}}$, this gives us the resulting regret bound. In the simple case of a complete network on N nodes, no coordinator is required, and we obtain the following corollary.

Corollary 4.7. *On a complete network with b nodes, if all nodes use the* EXPN *algorithm, then they attain $R = O\left(\sqrt{\left(1+\frac{K}{b}\right) T \ln K}\right)$, which is optimal (up to log factors) for any centralized solution.*

This again follows from the proof of Theorem 4.1 using the fact that the number of neighbors is $b-1$ on a complete network, and that a centralized solution has average regret $\Omega(\sqrt{\left(1+\frac{K}{b}\right) T})$ as shown in [AKTT15].

5 Empirical Results

5.1 Adversarial Setting

Benchmarks. We compare our algorithm against the bandit algorithms developed for various settings with side-information, namely EXP3G ([ACBDK15]), EXP.IX ([KNVM14]) and BEXP ([AKTT15]). The first two are designed for the arm-network setting as described in Section 4.5, while the latter is designed for the free-exploration setting described in Section 2. Recall that in free-exploration there are no neighbors; rather there is a budget B, and at each time step the individual can choose up to B arms to select. In order to attain a fair comparison, we assume we have budget $B = b + 1$ for BEXP, where b is the number of neighbors.

Experimental Setup. For the simulations we use the decoupled version of the EXPN algorithm as presented in Section 4.3; the results for the adaptive algorithm version would be even better. We consider a bandit with Bernoulli rewards that has a single *good* arm with mean 0.7, while the remaining arms have mean 0.5. This is similar to the worst case (minimax) bandit; the difficulty arises from the fact that it is hard, in an information theoretic sense, to distinguish the single good arm from the rest with few samples. Indeed the performance for our algorithm in comparison to our benchmarks is only improved for all other settings we attempted.

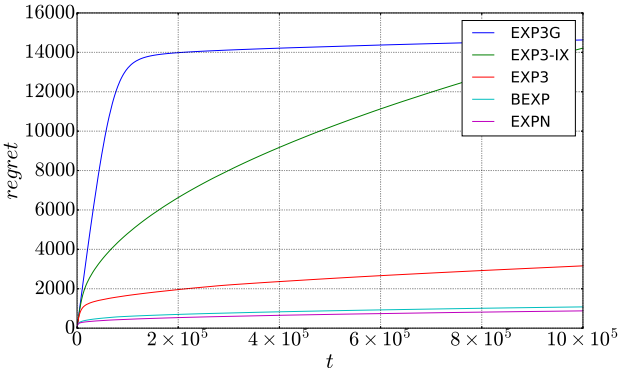

(a) Regret in a complete network on 5 nodes. We vary T for $K = 50$.

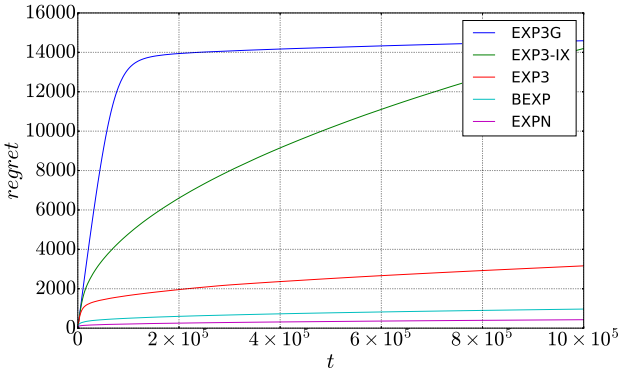

(b) Regret in a random 5-regular network with 50 nodes. We vary T for $K = 50$.

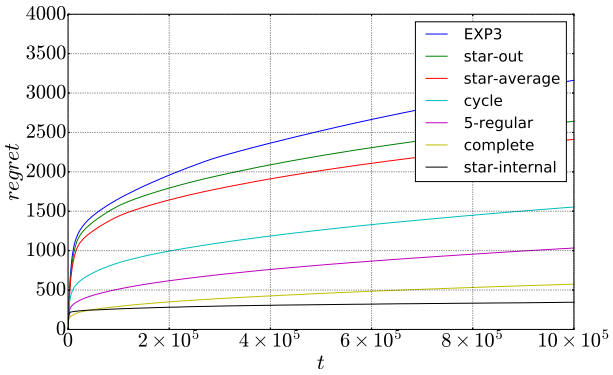

(c) The average regret in various network topologies. The size of the networks is 10.

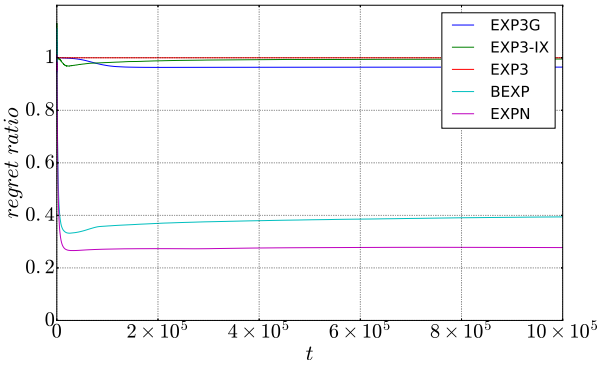

(d) Regret ratio. We vary time T for $K = 50$ and $N = 5$.

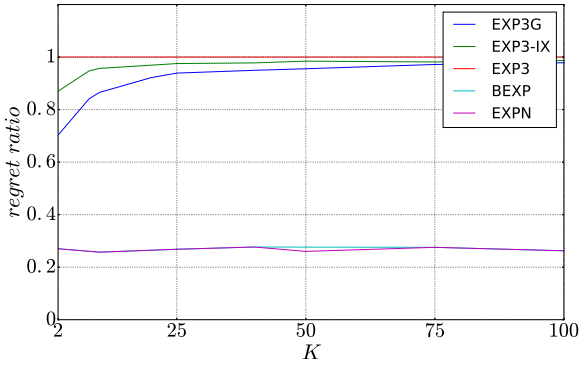

(e) Regret ratio. We vary the number of arms K for $T = 5 \cdot 10^4$ and $N = 5$.

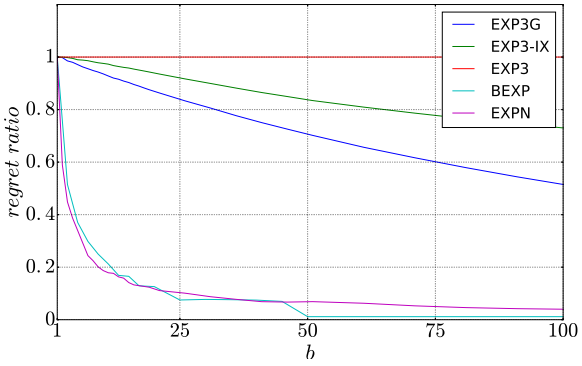

(f) Regret ratio. We vary the number of nodes N for $K = 50$ and $T = 10^6$.

Figure 3: Performance of our algorithm (EXPN) for the adversarial setting against benchmarks. Our algorithm significantly outperforms EXP3, indicating that the presence of neighbors indeed improves learning. It also significantly outperforms arm-network algorithms (EXP3G and EXP.IX) that could be applied in our setting. Surprisingly, it's performance is as good as BEXP, which would require a single node to dictate the choices of her neighbors; hence, our distributed algorithm is performing as well as a centralized one. Figures (a)-(c) depict the regret. Figures (d)-(f) depict the regret ratio, i.e., the ratio between an algorithm's regret with b neighbors over its regret with 0 neighbors (where b depends on the network structure addressed in the corresponding subfigure).

Performance in Networks. In addition to exploring the effect of the various algorithm on a single individual, we are able to consider various network topologies and consider the regret as a whole. Towards this, in the first set of simulations, all nodes in the specified networks use the same algorithm. We first compare the regret of EXPN against the benchmarks in a complete network on 5 nodes (Figure 3(a)); even on such a small network the difference in regret is dramatic.[6] We significantly outperform arm-network algorithms (EXP3G and EXP.IX), which empirically are initially worse than even EXP3. Asymptotically EXP3G eventually outperforms EXP3, although EXP.IX does not. Surprisingly, our algorithm performs as well as BEXP, which would be equivalent to identifying a single node as the leader and having them dictate the action of all other nodes. Hence, our distributed algorithm is as good as a centralized one. For comparison, we also consider a random 5-regular graph on 50 nodes (Figure 3(b)), and observe that the performance of all algorithms is roughly equivalent to the complete network on 5 nodes; i.e., the primary determining factor in the regret appears to be the number of neighbors rather than the topology of the network.

We also consider the regret of EXPN on various network topologies on 10 vertices: the complete network, a random 5-regular network, a cycle, and a star network (Figure 3(c)). When the number of neighbors differ in a topology, the regret of the nodes may differ; the star is the extreme example and we depict the minimum (for the center node), maximum (for one of the leaves) and average regret. As expected, the more neighbors one has, the better the regret is, with the internal node of star outperforming all. We also observe that there is an advantage to having neighbors that are not well-connected; despite a node in the complete network having the same degree as the center node of the star, the former has more regret. Because the nodes that are not well-connected receive less information, they must explore more – this is advantageous for their neighbors.

Performance of Individuals. Moving back to analyzing the performance for an individual, consider a setting where her neighbors all use the EXP3 algorithm. We measure the *regret ratio*, i.e., the ratio between the regret of bandit algorithm $\mathcal{A}$ when the node has b neighbors divided by the regret of $\mathcal{A}$ when the node has 0 neighbors. This allows us to better visualize the improvement in regret that each algorithm obtains as a function of the number of neighbors. We vary time T (Figure 3(d)), the number of arms K (Figure 3(e) and the number of neighbors b (Figure 3(f)). We observe that, in all cases, our EXPN algorithm always matches or outperforms the benchmarks. The fact that the performance of our EXPN is comparable to that of BEXP is surprising, as we could not hope to do any better.

5.2 Stochastic Setting

The setup for the empirical results in this section parallel that of Section 5.1. Recall that we make no assumption in our algorithm about our neighbors or how they play. We simply observe their actions and rewards. We let $\alpha = 2.5$ in the UCBN algorithm; the performance could be improved by optimizing α. We first observe that more neighbors leads to less regret (Figure 4(f)).

We then consider the regret of UCBN on various network topologies on 10 vertices: the complete network, a random 5-regular network, and a star network (Figure 4(c)). Similar to the previous experiments, for networks in which all vertices have the same number of neighbors (all but the star network), all agents attain the same regret and hence we report the average regret. However, this is not the case if the number of neighbors differ; the star is the extreme example and we depict the minimum (for the center node), maximum (for one of the leaves) and average regret. As expected, the more neighbors one has, the better the regret is, with the complete network and center of star outperforming all. We also observe that there is advantage to having neighbors that are not well-connected; despite a node in the complete network having the same degree as the center node of the star, the former has slightly more regret. The reason is a neighbor with lower degree attains less information from neighbors and explore the suboptimal arms more which itself is in the favor of its neighbors (here the center of star).

We then consider the *regret ratio*, i.e., the ratio between the regret of bandit algorithm $\mathcal{A}$ when the agent has b neighbors divided by the regret of $\mathcal{A}$ when the agent has 0 neighbors. This allows us to better visualize the improvement in regret that each algorithm obtains as a function of the number of neighbors. We vary the number of arms K (Figure 4(e) and the number of neighbors b (Figure 4(f)). We observe that, in all cases, our algorithm UCBN attains the theoretical regret ratio, i.e, in the complete graph when all agents use UCBN the regret ration $\mathcal{A}_b \to 1/b$ as $T \to \infty$.

[6] Indeed, on larger networks the differences are only more pronounced – we present the results on a small network in order to be able to visualize them adequately.

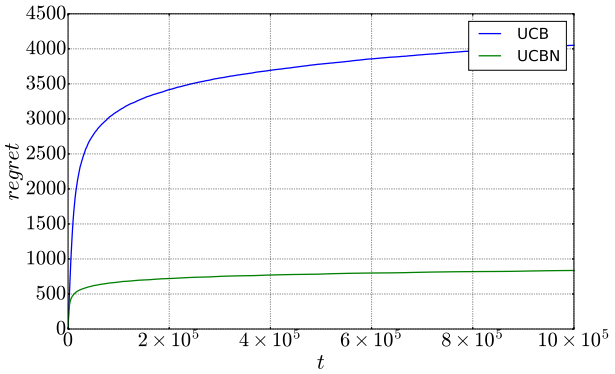

(a) Regret in a complete network on 5 agents. We vary T for $K = 5$.

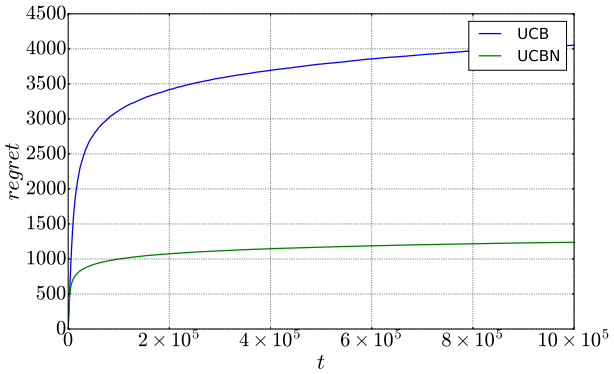

(b) Regret in a random 5-regular network with 10 agents. We vary T for $K = 5$.

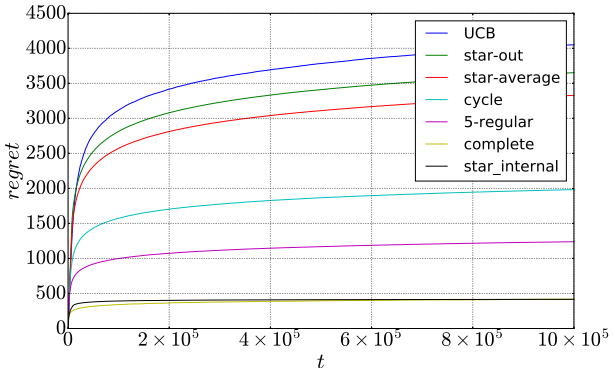

(c) The average regret in various network topologies. The size of the networks is 10, $K = 50$ and all agents use UCBN.

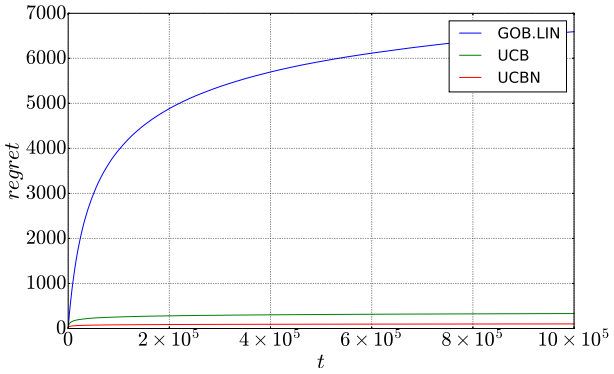

(d) Regret in the complete network when we vary the number of agents for $K = 5$ and $T = 25 \cdot 10^4$.

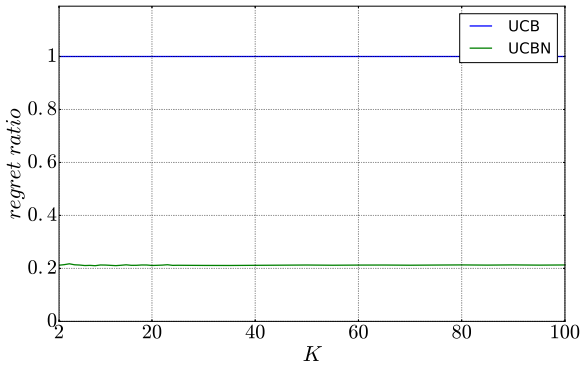

(e) Regret ratio in the complete network on 5 agents. We vary K for $T = 25 \cdot 10^4$.

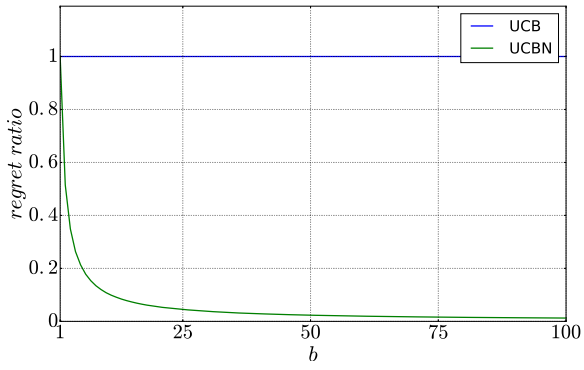

(f) Regret ratio in the complete network when we vary the number of agents for $K = 50$ and $T = 25 \cdot 10^4$.

Figure 4: Performance of our algorithm (UCBN) for the stochastic setting against benchmarks. Our algorithm significantly outperforms UCB, indicating that the presence of neighbors indeed improves learning. It also significantly outperforms GOB.LIN that could be applied in our setting (Figure 4(d)). Figures (a)-(d) depict the regret. Figures (e)-(f) depict the regret ratio, i.e., the ratio between an algorithm's regret with b neighbors over its regret with 0 neighbors (where b depends on the network structure addressed in the corresponding subfigure).

Finally, we compare our algorithm to the one proposed in [CBGZ13] (GOB.LIN); see Figure 4(d). Although this algorithm is a centralized and developed for a different setting (namely, for linear contextual bandits), it can be adapted to our setting by assuming that individuals are cooperative instead of selfish. Despite the centralized nature of GOB.LIN, our algorithm outperforms its regret.

6 Conclusion & Future Work

In this paper, we consider a model for social learning that puts the problem in the the bandit framework This model allows the problem to be analyzed both in the stochastic and adversarial bandit setting, and we provide algorithms for both cases. The regret of our algorithms interpolates between the regret of the traditional bandit setting (e.g., when an individual has no neighbors) and the regret of the full information setting (e.g., when the number of neighbors goes to infinity), and are optimal up to log factors. We show, both theoretically and empirically, that we outperform state-of-the-art bandit algorithms that one could also apply to this setting, and illustrate how our approach could also lead to centralized algorithms of interest.

With respect to improvements to the social learning model, relaxing assumption (3) would be ideal. As we have shown (see Proposition 4.5), removing it entirely results in strictly weaker regret bounds. Would an alternate relaxed assumption suffice? Lastly, it remains to formally study the effect of arbitrary network topologies on the regret, both for the individual (based on their position in the network) and on average.

References

[ACBDK15] N. Alon, N. Cesa-Bianchi, O. Dekel, and T. Koren. Online learning with feedback graphs: Beyond bandits. In *Conference on Learning Theory (COLT)*, 2015.

[ACBF02] P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47:235–256, 2002.

[ACBFS02] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire. The non-stochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32:47–77, 2002.

[ACBFS03] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. Schapire. The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 2003.

[ACBG02] Peter Auer, Nicolo Cesa-Bianchi, and Claudio Gentile. Adaptive and self-confident on-line learning algorithms. *Journal of Computer and System Sciences*, 64(1):48–75, 2002.

[ACBG$^+$14] N. Alon, N. Cesa-Bianchi, C. Gentile, S. Mannor, and Y. Mansour. Nonstochastic multi-armed bandits with graph-structured feedback. *Arxiv*, 2014.

[ACBGM13] N. Alon, N. Cesa-Bianchi, C. Gentile, and Y. Mansour. From bandits to experts: A tale of domination and independence. In *Proceedings of the 26th Conference on Advances in Neural Information Processing Systems (NIPS)*, 2013.

[AG12] Shipra Agrawal and Navin Goyal. Analysis of thompson sampling for the multi-armed bandit problem. In *COLT*, pages 39–1, 2012.

[AHK12] S. Arora, E. Hazan, and S. Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8:121–164, 2012.

[AKTT15] Kareem Amin, Satyen Kale, Gerald Tesauro, and Deepak Turaga. Budgeted prediction with expert advice. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.

[AMS12] O. Avener, S. Mannor, and O. Shamir. Decoupling exploration and exploitation in multi-armed bandits. In *Proceedings of the 29th International Conference on Machine Learning (ICML)*, 2012.

[ASC12] E. Accinelli and J. Sánchez-Carrera. Corruption driven by imitative behavior. *Econ. Letters*, 117:84–87, 2012.

[BCB12] S. Bubeck and N. Cesa-Bianchi. *Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems*, volume 5. Foundations and Trends in Machine Learning, 2012.

[BES13] S. Buccapatnam, A. Eryilmaz, and N. B. Shroff. Multi-armed bandits in the presence of side observations in networks. In *Proceedings of the 2014 SIGMETRICS conference*, 2013.

[BES14] S. Buccapatnam, A. Eryilmaz, and N. B. Shroff. Stochastic bandits with side observations on networks. In *Proceedings of the 52nd IEEE Conference on Decision and Control*, 2014.

[BG98] V. Bala and S. Goyal. Learning from neighbours. *Review of Econ. Studies*, 65:595–621, 1998.

[BG01] V. Bala and S. Goyal. Conformism and diversity under social learning. *Econ. Theory*, 17:101–120, 2001.

[Bub10] S. Bubeck. *Bandits games and clustering foundations*. PhD thesis, Universite Lille, 2010.

[CBGZ13] N. Cesa-Bianchi, C. Gentile, and G. Zappella. A gang of bandits. In *Proceedings of the 26th Conference on Advances in Neural Information Processing Systems (NIPS)*, 2013.

[CKLB12] S. Caron, B. Kveton, M. Lelarge, and S. Bhagat. Leveraging side observations in stochastic bandits. In *Proceedings of Uncertainty in Artificial Intelligence (UAI)*, 2012.

[EF93] G. Ellison and D. Fudenberg. Rules of thumb for social learning. *J. of Political Economy*, 101, 1993.

[FKM05] A. D. Flaxman, A. T. Kalai, and H. B. McMahan. Online convex optimization in the bandit setting: gradient descent without a gradient. In *Proceedings of the 16th ACM/SIAM symposium on Discrete algorithms (SODA)*, 2005.

[GC11] A. Garivier and O. Cappé. The kl-ucb algorithm for bounded stochastic bandits and beyond. In *Proceedings of the Conference on Learning Theory (COLT)*, 2011.

[GJ10] B. Golub and M. O. Jackson. Naive learning in social networks and the wisdom of crowds. *American Econ. Journal: MicroEcon.*, 2, 2010.

[GK03] D. Gale and S. Kariv. Bayesian learning in social networks. *Games and Econ. Behavior*, 45, 2003.

[Goy05] S. Goyal. Learning in networks. In G. Demange and M. Wooders, editors, *Group formation in Econ.: networks, clubs, and coalitions*, chapter 4, pages 122–167. Cambridge University Press, 2005.

[KL55] E. Katz and P. Lazersfeld. *Personal Influence*. The Free Press, 1955.

[KNVM14] T. Kocák, G. Neu, M. Valko, and R. Munos. Efficient learning by implicit exploration in bandit problems with side observations. In *Proceedings of the 27th Conference on Advances in Neural Information Processing Systems (NIPS)*, 2014.

[LBG48] P. Lazarsfeld, B. Berelson, and H. Gaudet. *The People's Choice*. Columbia University Press, 1948.

[LR85] T. L. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6, 1985.

[MMS11] O. A. Maillard, R. Munos, and G. Stoltz. A finite-time analysis of multiarmed bandits problems with kullback-leibler divergences. In *Proceedings of the Conference on Learning Theory (COLT)*, 2011.

[MS11] S. Mannor and O. Shamir. From bandits to experts: On the value of side-observations. In *Proceedings of Advances in Neural Information Processing Systems (NIPS)*, pages 684–692, 2011.

[San06] B. Sanditov. *Essays on Social Learning and Imitation*. PhD thesis, Universitaire Pers Maastricht, 2006.

[SBFH+13] Balazs Szorenyi, Róbert Busa-Fekete, István Hegedüs, Róbert Ormándi, Márk Jelasity, and Balázs Kégl. Gossip-based distributed stochastic bandit algorithms. In *30th International Conference on Machine Learning (ICML 2013)*, volume 28, pages 19–27. Acm Press, 2013.

[YM09] J. Y. Yu and S. Mannor. Piecewise-stationary bandit problems with side observations. In *Proceedings of the 26th International Conference on Machine Learning (ICML)*, 2009.

[Yoo12] D. Yoo. Individual and social learning in bio-technology adoption: The case of gm corn in the u.s. In *Agricultural & Applied Econ. Association's Annual Meeting (AAEA)*, 2012.

[ZAA07] J. Zhang, M. S. Ackerman, and L. Adamic. Expertise networks in online communities: Structures and algorithms. In *Proceedings of the World Wide Web Conference (WWW)*, 2007.

A Adversarial Bandits

A.1 Regret bound for Equation 4

We first state and prove a slightly simpler regret upper bound that is more intuitive, and then show how to enhance the proof to give Theorem 4.1 in Appendix A.2. In this regret bound the algorithm is as described in the main body of the paper, and we use a fixed (as opposed to adaptive) parameters δ and η.

Theorem A.1. *Given an agent with b neighbors who are playing arbitrarily, the regret of an agent using the EXPN algorithm is*

$$R \in O(\sqrt{\beta T \ln K})$$

where

$$\Theta = \Pi_{i=1}^{b}\left(1-\frac{\varepsilon_i}{K}\right) \text{ and } \beta = \frac{1}{1-\left(1-\frac{1}{K}\right)\Theta}+1.$$

for an optimal choice of η and δ that depends on Θ.

More precisely, we show that

$$R \leq \begin{cases} 2\sqrt{\beta T \ln K} & \text{if } \Theta \leq 1-\sqrt{\frac{\ln K}{\beta T}}, \text{ for } \eta=0, \delta=\sqrt{\frac{\ln K}{\beta T}}, \\ \sqrt{\beta T \ln K}+\sqrt{(\beta+\frac{K}{\Theta})T\ln K} & \text{if } 1-\sqrt{\frac{\ln K}{\beta T}} \leq \Theta \leq 1-\sqrt{\frac{\ln K}{(\beta+\frac{K}{\Theta})T}}, \text{ for } \eta=0, \delta=1-\Theta, \\ 2\sqrt{\left(\beta+\frac{K}{\Theta}\right)T\ln K} & \text{if } 1-\sqrt{\frac{\ln K}{(\beta+\frac{K}{\Theta})T}} \leq \Theta \\ & \text{for } \eta=\frac{K}{\Theta}\left(\sqrt{\frac{\ln K}{(\beta+\frac{K}{\Theta})T}}+(\Theta-1)\right), \delta=\sqrt{\frac{\ln K}{(\beta+\frac{K}{\Theta})T}}. \end{cases}$$

We can now directly reinterpret this regret (R_{EXPN}) as a function of the bandit regret (R_{EXP3}) and full information information regret (R_{MUA}) as in Equation 4. In particular, note that when $\Theta = 1$, none of our neighbors maintain a probability distribution that is bounded away from 0 for all arms. Hence, our neighbors are effectively not exploring. In this case, $\beta = K+1$ and $R_{\mathrm{EXPN}} \in O(\sqrt{TK\ln K})$, the same as in the classical bandit setting. On the other hand, when $\Theta \leq 1/2$, then $\beta \leq 3$ and hence $R_{\mathrm{EXPN}} \in O(\sqrt{T\ln K})$, the same as in the full-information setting. Hence, this algorithm smoothly interpolates between the bandit regret to full information regret as a function of the neighbor's exploration.

Proof. First, note that $\mathbb{E}[\widehat{g}_j(t)] = \left(1-p'_j(t)\right)\cdot 0 + p'_j(t)\frac{g_j(t)}{p'_j(t)} = g_j(t)$; hence we are indeed using an unbiased estimator for the rewards. From the definition of $w_j(t)$ (see Section 4.3), we see that

$$\ln\frac{W_{T+1}}{W_0} \geq \ln\frac{w_j(T+1)}{W_0} \geq \delta\sum_{t=1}^{T}\widehat{g}_j(t) - \ln K. \tag{16}$$

Moreover, using the definition of $p_j(t)$ (see Section 4.3),

$$\frac{W_{t+1}}{W_t} = \sum_{j=1}^{K} \frac{w_j(t)}{W_t} e^{\delta \widehat{g}_j(t)} = \sum_{j=1}^{K} \left(\frac{p_j(t) - \frac{\eta}{K}}{1-\eta} \right) e^{\delta \widehat{g}_j(t)}.$$

Since the algorithm selects a δ such that $\delta \widehat{g}_j(t) < 1$, we can use the inequality $e^z \leq 1 + z + z^2$, which holds for all $z \leq 1$. Therefore,

$$\begin{aligned} \frac{W_{t+1}}{W_t} &\leq \sum_{j=1}^{K} \left(\frac{p_j(t) - \frac{\eta}{K}}{1-\eta} \right) (1 + \delta \widehat{g}_j(t) + (\delta \widehat{g}_j(t))^2) \\ &\leq 1 + \frac{\delta}{1-\eta} \sum_{j=1}^{K} p_j(t) \widehat{g_j} + \frac{\delta^2}{1-\eta} \sum_{j=1}^{K} p_j(t) \widehat{g_j}^2. \end{aligned} \tag{17}$$

Taking logarithms of Equation (17) and using the inequality $\ln(1+x) < x$ which holds for all $x > 0$, we get

$$\ln \frac{W_{t+1}}{W_t} \leq \frac{\delta}{1-\eta} \sum_{j=1}^{K} p_j(t) \widehat{g_j} + \frac{\delta^2}{1-\eta} \sum_{j=1}^{K} p_j(t) \widehat{g_j}^2. \tag{18}$$

Since

$$\ln \frac{W_{T+1}}{W_0} = \sum_{t=0}^{T} \ln \frac{W_{t+1}}{W_t},$$

combining Equations (16) and (18) with above and noting that $\widehat{g}_i(0) = 0$ we have that

$$\delta \sum_{t=1}^{T} \widehat{g}_j(t) - \ln K \leq \frac{\delta}{1-\eta} \sum_{t=1}^{T} \sum_{j=1}^{K} p_j(t) \widehat{g_j}(t) + \frac{\delta^2}{1-\eta} \sum_{t=1}^{T} \sum_{j=1}^{K} p_j(t) (\widehat{g_j}(t))^2. \tag{19}$$

Now, we note that given $p_j(t)$ at time t we have

$$\mathbb{E}\left[\widehat{g_j}(t)\right] = g_j(t), \tag{20a}$$

$$\mathbb{E}\left[\sum_{t=1}^{T} \sum_{i=1}^{K} p_j(t) \widehat{g}_j(t) \right] = \sum_{t=1}^{T} \sum_{j=1}^{K} p_j(t) g_j(t) = \mathbb{E}\left[\sum_{t=1}^{T} g_{a(t)}(t) \right], \text{ and} \tag{20b}$$

$$\mathbb{E}\left[\sum_{t=1}^{T} \sum_{j=1}^{K} p_j(t) \widehat{g}_j^2(t) \right] = \sum_{t=1}^{T} \sum_{j=1}^{K} \frac{p_j(t)}{p'_j(t)} g_j^2(t). \tag{20c}$$

Where the expectation is over randomness of the algorithm. We will upper bound Equation (49c). Recall that $g_j(t) \leq 1$ for all j, t. Hence, for all t,

$$\sum_{j=1}^{K} \frac{p_j(t)}{p'_j(t)} g_j^2(t) \leq \max \left\{ \sum_{j=1}^{K} \frac{p_j(t)}{p'_j(t)} \right\}.$$

Where the maximization is over the space of valid probabilities for actions. This is upper bounded by Lemma A.2

$$\beta \stackrel{\text{def}}{=} \frac{1}{1 - (1 - \frac{1}{K})\Theta} + 1.$$

Since Equation (19) holds for all j, by (49), and Lemma A.2, we get

$$\delta \max_j \left[\sum_{t=1}^{T} g_j(t)\right] - \ln K \leq \frac{\delta}{1-\eta}\mathbb{E}\left[\sum_{t=1}^{T} g_{a(t)}(t)\right] + \frac{\delta^2}{1-\eta}\beta T.$$

Since $g_i(t) \leq 1$, we rearrange to get the following upper bound on the regret of our algorithm

$$R \leq \frac{\ln K}{\delta} + \eta T + \delta\beta T.$$

What remains is then an optimization problem in δ and η which is subject to the following two constraints:

$$\eta \in [0,1] \quad \text{and} \quad \frac{\delta}{1-(1-\frac{\eta}{K})\Theta} \in [0,1]. \tag{22}$$

In this optimization problem the only assumption made on the algorithm of agents is that they are select arms randomly and that the probability of selecting an arm has a minimum value ε.

$$\begin{aligned} &\min_{\delta,\eta} f(\delta,\eta) = \frac{\ln K}{\delta} + \eta \cdot T + \delta\beta T \\ &\text{Subject to} \\ &\qquad g_1(\delta,\eta) \leq 0, \\ &\qquad g_2(\delta,\eta) \leq 0, \end{aligned} \tag{23}$$

where

$$\begin{aligned} g_1(\delta,\eta) &= -\eta, \\ g_2(\delta,\eta) &= \frac{\delta}{1-\Theta(1-\frac{\eta}{N})} - 1, \end{aligned} \tag{24}$$

and

$$\begin{aligned} \Theta &= \Pi_{i=2}^{b}(1-\frac{\eta_i}{K}), \\ \beta &= \frac{1}{1-(1-\frac{1}{K})\Theta} + 1, \end{aligned} \tag{25}$$

If $\delta^\star$ and $\eta^\star$ is a local minimum that satisfies Karush-Kuhn-Tucker (KKT) conditions(see below).
Stationary:

$$-\nabla f(\delta^\star,\eta^\star) = \mu_1 \cdot \nabla g_1(\delta^\star,\eta^\star) + \mu_2 \cdot \nabla g_2(\delta^\star,\eta^\star), \tag{26}$$

Primal feasibility:

$$\begin{aligned} g_1(\delta^\star,\eta^\star) &\leq 0 \\ g_2(\delta^\star,\eta^\star) &\leq 0, \end{aligned} \tag{27}$$

Dual feasibility:

$$\begin{aligned} \mu_1 &\geq 0 \\ \mu_2 &\geq 0 \end{aligned} \tag{28}$$

Complementary slackness:

$$\begin{aligned} \mu_1 \cdot g_1(\delta^\star,\eta^\star) &= 0 \\ \mu_2 \cdot g_2(\delta^\star,\eta^\star) &= 0 \end{aligned} \tag{29}$$

In each step we will assume that some of these constraints are active (i.e.,$g_i(\delta,\eta) = 0$), and find the points that satisfy the KKT conditions,

$$-\nabla f(\delta,\eta) = \begin{pmatrix} \frac{\ln K}{\delta^2} - T\beta \\ -T \end{pmatrix} \tag{30}$$

First, let us assume that only first constraint is active,

$$g_1(\delta, \eta) = 0.$$

Which yields that

$$\eta^\star = 0.$$

The only stationary point in this case is

$$\delta^\star = \sqrt{\frac{\ln K}{\beta T}}. \tag{31}$$

And the Θ that primal and dual feasibility holds for is

$$\Theta \leq 1 - \sqrt{\frac{\ln K}{\beta T}}. \tag{32}$$

Since $\delta^\star$ and $\eta^\star$ satisfy KKT conditions they are valid answers for this interval. The regret is

$$R \leq \sqrt{\beta \ln KT}. \tag{33}$$

Second, let us assume that only the second constraint is active. In this case the stationary point is

$$\begin{aligned} \delta^\star &= 1 - \Theta(1 - \frac{\eta}{K}) \\ \eta^\star &= \frac{K}{\Theta}\sqrt{\frac{\ln K}{(\beta + K/\Theta)T}} + \frac{K(\Theta - 1)}{\Theta}. \end{aligned} \tag{34}$$

The interval that this answer is valid for is as follows,

$$\Theta \geq 1 - \sqrt{\frac{\ln K}{(\beta + \frac{K}{\Theta})T}}. \tag{35}$$

The regret is

$$R \leq 2\sqrt{(\beta + \frac{K}{\Theta})T \ln K} + \frac{K(\Theta - 1)T}{\Theta}. \tag{36}$$

Finally, let us assume that all of the constraints are active. In this case the stationary point is

$$\begin{aligned} \delta^\star &= 1 - \Theta \\ \eta^\star &= 0\,. \end{aligned} \tag{37}$$

And the interval that this answer is valid is as follows,

$$\begin{aligned} \Theta &\geq 1 - \sqrt{\frac{\ln K}{\beta T}}, \\ \Theta &\leq 1 - \sqrt{\frac{\ln K}{(\beta + \frac{K}{\Theta})T}}\,. \end{aligned} \tag{38}$$

The regret is

$$R \leq \frac{\ln K}{1 - \Theta} + (1 - \Theta)\beta T\,. \tag{39}$$

After solving the inequalities (32), (35) and (38) we get the desired bound for the regret. □

Lemma A.2. $\sum_{j=1}^{K} \frac{p_j(t)}{p'_j(t)} \leq \frac{1}{1-(1-\frac{1}{K})\Theta} + 1 \stackrel{\text{def}}{=} \beta$ *for all* t.

Proof of Lemma A.2. By definition (see Equations 5 and 6),

$$S \stackrel{\text{def}}{=} \sum_{j=1}^{K} \frac{p_j(t)}{p'_j(t)} = \sum_{j=1}^{K} \frac{p_j(t)}{1-(1-p_j(t))\cdots(1-q_j^b(t))}.$$

Since $q_j^i \geq \varepsilon_i/K$, clearly

$$S \leq \sum_{j=1}^{K} \frac{p_j(t)}{1-(1-p_j(t))\Theta}.$$

Notice that the right-hand side is strictly concave with respect to each p_i. Thus, the maximum is achieved either on the boundary of the feasible region of p, or at a single point in the interior. Recall that $p_i \in [\eta/K, 1-(K-1)\frac{\eta}{K}]$, and that p is a probability distribution. Hence, boundary points are of the form η/K for all except one entry, which is $1-(K-1)\frac{\eta}{K}$. If p is of this form, then

$$S \leq \frac{(K-1)}{K} \cdot \frac{\eta}{1-(1-\frac{\eta}{K})\Theta} + \frac{1-\frac{(K-1)\eta}{K}}{1-\frac{(K-1)\eta\Theta}{K}} \leq \frac{\eta}{1-(1-\frac{\eta}{K})\Theta} + 1. \tag{40}$$

If, instead, the maximum is achieved on the interior, it must be symmetric, i.e., the probability of playing all actions is $1/K$; otherwise, since the equation is symmetric, there would be more than one maximal distribution which contradicts strict concavity. If p is of this form, then

$$S \leq \frac{1}{1-(1-\frac{1}{K})\Theta}. \tag{41}$$

Since $\eta \in [0,1]$, we can upper bound both the right-hand sides of Equations 40 and 41 to attain the desired bound:

$$\sum_{i=1}^{K} \frac{p_i(t)}{p'_i(t)} \leq \frac{1}{1-(1-\frac{1}{K})\Theta} + 1.$$

□

A.2 Proof of Main Theorem 4.1

For this result, we use the same unbiased estimator as presented in the main body of the paper, but take a different approach in the analysis. Instead of having an explicit exploration parameter $\eta > 0$, we instead will allow the update parameter δ_t to vary with time t (in fact, in this case $\eta = 0$). This sort of analysis is common (see, e.g., [BCB12] and [KNVM14]), in particular in settings where T is unknown. In our case the adaptive δ_t functions to incorporate information from neighbors as we see it.

Proof of Main Theorem 4.1. The first part of proof (from Equation (42) to Equation (48)) parallels the format of the proof of Theorem 3.1 in [BCB12] for the usual bandit setting. We first write the $\mathbb{E}[\widehat{l}_j(t)]$ as Equation (42), then we upper bound the first term and lower bound the second term, which leads to Equation (48).

$$\mathbb{E}[\widehat{l}_j(t)] = \frac{1}{\delta}\left(\ln \mathbb{E}\left[\exp\left(-\delta\left(\widehat{l}_j(t) - \mathbb{E}[\widehat{l}_j(t)]\right)\right)\right] - \ln \mathbb{E}\left[\exp\left(-\delta \widehat{l}_j(t)\right)\right]\right). \tag{42}$$

where the expectation is over the randomness of the estimator and choice of the arm.

We now diverge from the usual proof template: In the next step we find an upper bound for the first term in right-hand side of above equation.

$$\begin{aligned}\frac{1}{\delta}\ln\mathbb{E}\left[\exp\left(-\delta(\widehat{l}_j(t)-\mathbb{E}[\widehat{l}_j(t)])\right)\right] &= \frac{1}{\delta}\ln\mathbb{E}\left[\exp(-\delta\widehat{l}_j(t)\right]+\mathbb{E}[\widehat{l}_j(t)]) \\ &\leq \frac{1}{\delta}\mathbb{E}[\exp(-\delta\widehat{l}_j(t))-1+\delta\widehat{l}_j(t)] \\ &\leq \frac{\delta}{2}\mathbb{E}[\widehat{l}_j^2(t)],\end{aligned} \tag{43}$$

where in the second inequality we use $\ln x \leq x-1$ and in the last inequality we use $\exp(-x)-1+x \leq x^2/2$ for $x \geq 0$. Defining $\widehat{L}_j(t) = \sum_{a=1}^{t}\widehat{l}_j(t)$ we have

$$\begin{aligned}-\frac{1}{\delta}\ln\mathbb{E}_{a_t\sim p'_t}\mathbb{E}_{j\sim p'_t}\exp(-\delta\widehat{l}_j(t)) &\leq -\frac{1}{\delta}\mathbb{E}_{a_t\sim p'_t}\ln\mathbb{E}_{j\sim p'_t}\exp(-\delta\widehat{l}_j(t)) \\ &= -\frac{1}{\delta}\mathbb{E}_{a_t\sim p'_t}\ln\sum_{j=1}^{K}p'_j(t)\exp(-\delta\widehat{l}_j(t)) \\ &= -\frac{1}{\delta}\mathbb{E}_{a_t\sim p'_t}\left[\ln\left(\sum_{j=1}^{K}(1-\eta)\frac{\exp(-\delta\widehat{L}_j(t))}{\sum_{c=1}^{K}\exp(-\delta\widehat{L}_c(t-1))}+\frac{\eta}{K}\exp(-\delta\widehat{l}_j(t))\right)\right] \\ &\leq -\frac{1}{\delta}\mathbb{E}_{a_t\sim p'_t}\left(\frac{\eta}{K}\sum_{j=1}^{K}(-\delta\widehat{l}_j(t))+(1-\eta)\left[\ln\left(\sum_{j=1}^{K}\frac{\exp(-\delta\widehat{L}_j(t))}{\sum_{c=1}^{K}\exp(-\delta\widehat{L}_c(t-1))}\right)\right]\right),\end{aligned} \tag{44}$$

where in the first and last inequality we used the Jensen's inequality. Now we want to upper bound the term $-\frac{1}{\delta}\ln\left(\sum_{i=j}^{K}\frac{\exp(-\delta\widehat{L}_j(t))}{\sum_{c=1}^{K}\exp(-\delta\widehat{L}_c(t-1))}\right)$ in the above inequality.

$$-\frac{1}{\delta}\ln\left(\sum_{j=1}^{K}\frac{\exp(-\delta\widehat{L}_j(t))}{\sum_{c=1}^{K}\exp(-\delta\widehat{L}_c(t-1))}\right) = \psi(t-1)-\psi(t), \tag{45}$$

where $\psi(t) = \frac{1}{\delta}\ln\left(\frac{1}{K}\sum_{c=1}^{K}\exp(-\delta\widehat{L}_c(t))\right)$. By summing up these terms in Equation (42), (44), and (45) we get

$$\sum_{t=1}^{T}\mathbb{E}[\widehat{l}_j(t)] \leq \sum_{t=1}^{T}\frac{\delta}{2}\mathbb{E}[\widehat{l}_j^2(t)]+\sum_{t=1}^{T}\frac{\eta}{K}\mathbb{E}_{a_t\sim p'_t}\sum_{c=1}^{K}\widehat{l}_c(t)-\mathbb{E}_{a_t\sim p'_t}\psi(T). \tag{46}$$

By bounding $-\psi(T)$ we are able to find an upper bound on the regret.

$$\begin{aligned}-\psi(T) &= \frac{\ln K}{\delta}-\frac{1}{\delta}\ln\left(\sum_{j=1}^{K}\exp(-\delta\widehat{L}_j(T))\right) \\ &\leq \frac{\ln K}{\delta}-\frac{1}{\delta}\ln\left(\exp(-\delta\widehat{L}_k(T))\right) \\ &= \frac{\ln K}{\delta}+\widehat{L}_k(T).\end{aligned} \tag{47}$$

Plugging this in the above inequality yields

$$\sum_{t=1}^{T}\mathbb{E}[\widehat{l}_j(t)] \leq \sum_{t=1}^{T}\frac{\delta}{2}\mathbb{E}[\widehat{l}_j^2(t)]+\sum_{t=1}^{T}\frac{\eta}{K}\mathbb{E}_{a_t\sim p'_t}\sum_{c=1}^{K}\widehat{l}_c(t)+\frac{\ln K}{\delta}+\mathbb{E}_{a_t\sim p'_t}[\widehat{L}_k(T)]. \tag{48}$$

In fact, having more neighbors help us to get better bounds for $\mathbb{E}[\widehat{l}_j^2(t)]$.

Now, we note that given $p_j(t)$ at time t we have

$$\mathbb{E}\left[\widehat{g}_i(t)\right] = g_i(t), \tag{49a}$$

$$\mathbb{E}\left[\sum_{t=1}^{T}\sum_{i=1}^{K} p_i(t)\widehat{g}_i(t)\right] = \sum_{t=1}^{T}\sum_{i=1}^{K} p_i(t)g_i(t) = \mathbb{E}\left[\sum_{t=1}^{T} g_{a(t)}(t)\right], \text{ and} \tag{49b}$$

$$\mathbb{E}\left[\sum_{t=1}^{T}\sum_{i=1}^{K} p_i(t)\widehat{g}_i^2(t)\right] = \sum_{t=1}^{T}\sum_{i=1}^{K} \frac{p_i(t)}{p_i'(t)} g_i^2(t). \tag{49c}$$

Where the expectation is over randomness of the algorithm. Note that $g_i(t)$ is less than 1, as a result we have,

$$\sum_{t=1}^{T}\sum_{i=1}^{K} \frac{p_i(t)}{p_i'(t)} g_i^2(t) \leq \sum_{t=1}^{T}\sum_{i=1}^{K} \frac{p_i(t)}{p_i'(t)}.$$

Setting $\eta = 0$ and using the Lemma A.4 we have

$$R \leq \frac{\ln K}{\delta_T} + \frac{1}{2}\sum_{t=1}^{T} \delta_t(1+\gamma_t). \tag{50}$$

To conclude the proof, let $\delta_t = \sqrt{\frac{\ln K}{\sum_{c=1}^{t}(1+\gamma_t)}}$ and use Lemma 3.5 of [ACBG02], which gives an upper bound on the regret to get the expected regret, simply take expectation of the both sides of Equation (50).

□

Lemma A.3. *For $p^i \in [0,1]$ we want to show that $\Pi_{i=1}^{b}(1-p^i) \leq \frac{1}{1+\sum_{i=1}^{b} p^i}$.*

Proof. The following formula holds for all positive p_i,

$$\Pi_{i=1}^{b}(1+p^i) \geq 1 + \sum_{i=1}^{b} p^i, \tag{51}$$

which implies

$$\frac{1}{\Pi_{i=1}^{b}(1+p^i)} \leq \frac{1}{1+\sum_{i=1}^{b} p^i}. \tag{52}$$

As we have

$$\begin{aligned} &\Pi_{i=1}^{b}(1-p^i)\cdot\Pi_{i=1}^{b}(1+p^i) \\ &= \Pi_{i=1}^{b}(1-(p^i)^2) \leq 1, \end{aligned} \tag{53}$$

We can conclude

$$\Pi_{i=1}^{b}(1-p^i) \leq \frac{1}{\Pi_{i=1}^{b}(1+p^i)}. \tag{54}$$

From (52) and (54) we get

$$\begin{aligned} \Pi_{i=1}^{b}(1-p^i) &\leq \frac{1}{\Pi_{i=1}^{b}(1+p^i)} \\ &\leq \frac{1}{1+\sum_{i=1}^{b} p^i}. \end{aligned} \tag{55}$$

□

Lemma A.4. $\sum_{j=1}^{K} \frac{p_j}{1-(1-p_j)(1-q_j^1)\cdots(1-q_j^b)} \leq \sum_{j=1}^{K} \frac{p_j}{p_j+q_j^1+q_j^2+\cdots+q_j^b} + 1.$

Proof. From Lemma A.3 we have

$$(1-p_j)(1-q_j^1)\cdots(1-q_j^b) \leq \frac{1}{1+p_j+q_j^1+q_j^2+\cdots+q_j^b}, \tag{56}$$

substituting this bound in the statement of lemma yields

$$\begin{aligned}\sum_{j=1}^{K}\frac{p_j}{1-(1-p_j)(1-q_j^1)\cdots(1-q_j^b)} &\leq \sum_{j=1}^{K}\frac{p_j}{1-\frac{1}{1+p_j+q_j^1+q_j^2+\cdots+q_j^b}} \\ &\leq \sum_{j=1}^{K}\frac{p_j}{p_j+q_j^1+q_j^2+\cdots+q_j^b}+1,\end{aligned} \tag{57}$$

□

A.3 Proof of Regret Lower Bound

Proof of Theorem 4.2. This proof follows, with some adjustments, from the proof of Theorem 5.1 in [ACBFS03]; we simply point out the key differences here. In this proof a random distribution of rewards is constructed, one of the actions is chosen uniformly at random to be the "good" action, which is 1 with probability $1/2+\epsilon$ and 0 otherwise for some small fixed $\epsilon \in (0, 1/2]$. The remaining actions are 0 or 1 with probability $1/2$.

In Equation (30) of [ACBFS03], where the relative entropy between two Bernoulli random variables with parameters P and Q is calculated, the total number of past observation is T. In our case, it is between $\sum_{t=1}^{T} n_t$ and $T+\sum_{t=1}^{T} n_t$ where n_t is the size of the set of arms selected (arbitrarily) by all of the individual?s neighbors at time t. Note that, stopping here, this would lead directly to the proof of Theorem 4.5. Here, we dig into the proof to break up the regret in a way that allows us to write the regret with respect to our γ_t, which will result in a tighter bound.

More precisely, let r denote the entire sequence of rewards, let f is an arbitrary function on such histories r to a fixed range $[0, T]$, and let N_i be a random variable denoting the number of times an algorithm selects arm i. We let $\mathbb{E}_i$ be the expectation taken conditioned on i being the good action, and $\mathbb{E}_{unif}$ denotes the expectation taken over a uniformly random choice of rewards for all actions (including the good action). Then, Lemma **A.1** in [ACBFS03] becomes

$$\mathbb{E}_i[f(r)] \leq \mathbb{E}_{unif}[f(r)] + \frac{T}{2}\sqrt{-(\mathbb{E}_{unif}[N_i])\ln(1-4\epsilon^2)},$$

when $f(r)$ is the number of times that we choose arm i. Note that, for our algorithm, $\mathbb{E}_{unif}[N_i] = \sum_{t=1}^{T} 1-(1-p_i(t))(1-q_i^1(t))\cdots(1-q_i^b(t))$. Using the lemma, if we set $\epsilon = c\sqrt{\frac{K}{\sum_{i=1}^{K}\mathbb{E}_{unif}[N_i]}}$, and continue to follow the proof of Theorem 5.1 in [ACBFS03]. By observing that $\Omega\left(T\sqrt{\frac{K}{\sum_{i=1}^{K}\mathbb{E}_{unif}[N_i]}}\right) = \Omega\left(T+\sum_{t=1}^{T}\gamma_t\right)$, we recover the desired lower bound. □

A.4 Comparison to Arm-Network Algorithms

We compare our algorithm to arm-network algorithms which, though developed for a different setting, could be applied to ours. These algorithms are given for a multi-armed-bandit (MAB) with side observations encoded as an arm-network. In these papers, authors assume the arms form a arm-network G_t (the arm-network can change over time). The arm-network is a directed graph, whose vertices represent the arms, and an edge from arm i to arm j means that by choosing arm i we observe the reward of arm j (see Figure 2). First, let recall our interpretation of our problem as a multi-armed bandit problem with an arm-network (see Figure 2). The players choose arms according to their algorithm, and after selecting and revealing the rewards. Each player i individually can construct a arm-network. If we show the arms selected by neighbors of player i by $\mathcal{A}[N(i)]$, in arm-network there is an edge from all arms (vertices) to $\mathcal{A}[N(i)]$. In addition, if we denote the cardinality of set $\mathcal{A}[N(i)]$ by C_i, the independence number α of the underlying graph is $K+1-C_i$.

We show that our algorithm's regret is at most that of EXP3G, the state-of-the-art algorithm for the arm-network setting. From Theorem 4.1, we know that that the regret of our algorithm is $\widetilde{O}\left(\sqrt{T+\sum_{t=1}^{T}\gamma_t}\right)$ and from [ACBDK15] the regret of EXP3G is $\widetilde{O}\left(\sqrt{\sum_{t=1}^{T}\alpha_t}\right)$.

With some abuse of notation, let $\mathcal{A}$ be the arms chosen by neighbors of 0 let C be its cardinality, and let $1\{j\in\mathcal{A}\}$ be an indicator random variable that is 1 if and only if j is in $\mathcal{A}$. Lastly, let $1\{j=a_t^\ell\}$ be an indicator random variable that is 1 if and only if player ℓ (one of the neighbors) chooses arm j at round t. First, we lower bound α using the indicator random variable defined above, then we show that in expectation this term is greater than γ_t.

Lemma A.5. *The independence number α is lower bounded as follows*

$$\alpha_t \geq \sum_{j=1}^{K}\frac{p_j(t)}{p_j(t)+\sum_{\ell=1}^{b}1\{j=a_t^\ell\}}. \tag{58}$$

Proof. First, we show the following

$$\alpha_t \geq \sum_{j=1}^{K}\frac{p_j(t)}{p_j(t)+1\{j\in\mathcal{A}\}}. \tag{59}$$

Decomposing the sum we have

$$\sum_{j=1}^{K}\frac{p_j(t)}{p_j(t)+1\{j\in\mathcal{A}\}} = \sum_{j\in\mathcal{A}}\frac{p_j(t)}{p_j(t)+1}+\sum_{j\notin\mathcal{A}}\frac{p_j(t)}{p_j(t)}, \tag{60}$$

plugging the cardinality of $\mathcal{A}$ we get

$$= \sum_{j\in\mathcal{A}}\frac{p_j(t)}{p_j(t)+1}+K-C, \tag{61}$$

we know $\alpha = 1+K-C$ (this comes from the fact that arms in $\mathcal{A}$ are connected to all arms, see Figure 2), knowing that $\sum_{j\in\mathcal{A}}\frac{p_j(t)}{p_j(t)+1}$ is less than 1 completes the first part of lemma.

Second, we have $1\{j\in\mathcal{A}\}\leq\sum_{\ell=1}^{b}1\{j=a_t^\ell\}$, which yields the lemma. □

In the Lemma A.5, the term $1\{j=a_t^\ell\}$ can be seen as an unbiased estimator for $q_j^\ell(t)$ (we denote it by $\widehat{q}_j^\ell(t)$). As a final step, we show the following lemma.

Lemma A.6. *For a multinomial distribution q and their unbiased estimator $\widehat{q}_j^\ell(t)=1\{j=a_t^\ell\}$, we have*

$$\sqrt{\sum_{t=1}^{T}(1+\gamma_t)} \leq \mathbb{E}_{\widehat{q}\sim q}\left[\sqrt{2\sum_{t=1}^{T}\alpha_t}\right] \tag{62}$$

Proof. Because $\alpha_t\geq 1$, we know $\sum_{t=1}^{T}\alpha_t\geq T$. As a result by showing $\sqrt{\sum_{t=1}^{T}(\gamma_t)}\leq\mathbb{E}_{\widehat{q}\sim q}\left[\sqrt{\sum_{t=1}^{T}\alpha_t}\right]$, we can conclude $\sqrt{T+\sum_{t=1}^{T}(\gamma_t)}\leq\mathbb{E}_{\widehat{q}\sim q}\left[\sqrt{2\sum_{t=1}^{T}\alpha_t}\right]$.

From Lemma A.5, we have

$$\mathbb{E}_{\widehat{q}\sim q}\left[\sqrt{\sum_{t=1}^{T}\alpha_t}\right] \geq \mathbb{E}_{\widehat{q}\sim q}\left[\sqrt{\sum_{t=1}^{T}\sum_{j=1}^{K}\frac{p_j(t)}{p_j(t)+\sum_{\ell=1}^{b}\widehat{q}_j^\ell(t)}}\right]. \tag{63}$$

In the next step, we want to show

$$\mathbb{E}_{\widehat{q}\sim q}\left[\sqrt{\sum_{t=1}^{T}\sum_{j=1}^{K}\frac{p_j(t)}{p_j(t)+\sum_{\ell=1}^{b}\widehat{q}_j^{\ell}(t)}}\right]\geq\sqrt{\sum_{t=1}^{T}\gamma_t}.$$

Let $\phi(\widehat{q})=\phi(\widehat{q}^1(1),\widehat{q}^2(1),\cdots,\widehat{q}^b(1),\widehat{q}^1(2),\cdots,\widehat{q}^b(T))=\sqrt{\sum_{t=1}^{T}\sum_{j=1}^{K}\frac{p_j(t)}{p_j(t)+\sum_{\ell=1}^{b}\widehat{q}_j^{\ell}(t)}}$. This function is convex (it is convex along every arbitrary line with positive entires, so it is convex), we can use Jensen's inequality to swap the order of expectation and ϕ to get the following.

$$\mathbb{E}_{\widehat{q}\sim q}\left[\phi(\widehat{q})\right]\geq\phi\left(\mathbb{E}_{\widehat{q}\sim q}\left[\widehat{q}\right]\right)=\phi\left(q\right)=\sqrt{\sum_{t=1}^{T}\gamma_t}. \tag{64}$$

The first inequality is Jensen's inequality and the last equality comes from definition of γ_t. □

B Stochastic Bandits

In this section, we first formally state and prove the results for the stochatic setting. Let us first recall our result:

Theorem B.1 (Theorem 3.1)**.** *Consider an agent with neighbors who play arbitrarily. Let $n_i'(t)$ be the number times arm i has been selected by one of her neighbors by time t. Then, the regret of* UCBN *for any $\alpha>2$ is*

$$R\leq\sum_{i,\Delta_i>0}\left(\max\left\{\max_{t=1,..,T}\left\{\frac{2\alpha\ln t}{\Delta_i}-n_i'(t)\Delta_i\right\},0\right\}+\frac{\alpha}{\alpha-2}\right), \tag{65}$$

where Δ_i is the difference between $\mu_{i^\star}$ and μ_i.

Corollary B.2. *On a complete graph with b nodes, if all agents use* UCBN *then under the same conditions as in Theorem 3.1, the regret of an agent is*

$$R\leq\sum_{i,\Delta_i>0}\left(\frac{2\alpha\ln T}{b\Delta_i}+\frac{\alpha}{\alpha-2}\right)\in O\left(\frac{K\ln T}{b}\right). \tag{66}$$

The following lower bound yields same behavior for getting free observation from neighbors.

Theorem B.3. *Consider a strategy that satisfies $\mathbb{E}[n_i(T)]=o(T^a)$, any arm i with $\Delta_i>0$, and any $a>0$. Let $c_t(i)$ be the number of times arm i selected (arbitrarily) by all of the agent's neighbors up to time t, then, for any set of Bernoulli reward distributions the following inequality holds*

$$\lim_{T\longrightarrow+\infty}\inf\frac{R}{\ln T}\geq\sum_{i,\Delta_i>0}\frac{1}{2\Delta_i}-\lim_{T\longrightarrow+\infty}\inf\frac{\sum_{i,\Delta_i>0}c_T(i)\Delta_i}{\ln T}. \tag{67}$$

Our proofs parallel, with additional bookkeeping, the proofs for the original UCB results (see, e.g., [BCB12] for a template). Note that the results for stochastic bandits hold when the reward distributions satisfy the following standard conditions.

Definition B.1 (Conditions on $\mathcal{F}_i$)**.** *Every reward distribution $\mathcal{F}_i$ satisfies Hoefding's lemma, i.e., there exists a convex function ψ on the reals such that, for all $\lambda\geq0$, we have $\ln\left[\mathbb{E}\left[e^{\lambda|X-\mathbb{E}[X]|}\right]\right]\leq\psi(\lambda)$ where $X\sim\mathcal{F}_i$.*

For example, when $X\in[0,1]$, one can take $\psi(\lambda)=\frac{\lambda^2}{8}$; indeed the results in the main body of the paper take this ψ. The results can be easily generalized for other ψ in the usual manner.

We first prove a lemma that will be of assistance in the proof of Theorem 3.1. Recall that $a(t)$ is the arm the agent selects at time t.

Lemma B.4. *If $a(t) = i$, at least one of the three following inequalities is true:*

$$\widehat{\mu}_{i^\star, n_{i^\star}(t-1)} + \sqrt{\frac{\alpha \ln t}{2n_{i^\star}(t-1)}} \leq \mu^\star, \tag{68a}$$

$$\widehat{\mu}_{i, n_i(t-1)} > \mu_i + \sqrt{\frac{\alpha \ln t}{2n_i(t-1)}}, \tag{68b}$$

$$n_i(t-1) < \frac{2\alpha \ln t}{\triangle_i^2}. \tag{68c}$$

Proof. We prove the contrapositive. Assume that $a(t) = i$ and that none of the inequalities (68a), (68b) or (68c) are true.

$$\widehat{\mu}_{i^\star, n_{i^\star}(t-1)} + \sqrt{\frac{\alpha \ln t}{2n_{i^\star}(t-1)}} > \mu^\star, \tag{69a}$$

$$\widehat{\mu}_{i, n_i(t-1)} < \mu_i + \sqrt{\frac{\alpha \ln t}{2n_i(t-1)}}, \tag{69b}$$

$$n_i(t-1) > \frac{2\alpha \ln t}{\triangle_i^2}. \tag{69c}$$

By plugging $\mu^\star = \mu_i + \triangle_i$ in Equation (69a) we obtain

$$\widehat{\mu}_{i^\star, n_{i^\star}(t-1)} + \sqrt{\frac{\alpha \ln t}{2n_{i^\star}(t-1)}} > \mu_i + \triangle_i. \tag{70}$$

From Equation (69c) we have

$$\mu_i + \triangle_i > \mu_i + \sqrt{\frac{2\alpha \ln t}{n_i(t-1)}} \tag{71}$$

and plugging (71) in (70) yields

$$\widehat{\mu}_{i^\star, n_{i^\star}(t-1)} + \sqrt{\frac{\alpha \ln t}{2n_{i^\star}(t-1)}} > \widehat{\mu}_{i, n_i(t-1)} + \sqrt{\frac{\alpha \ln t}{2n_i(t-1)}}. \tag{72}$$

By our criteria for selecting arms given by Equation (1), this implies $a(t) \neq i$. □

Proof of Theorem 3.1 . Let

$$n_i(t) = n_i^1(t) + n_i'(t), \tag{73}$$

where $n_i^1(t)$ is the number of times the agent selects the arm i and $n_i'(t)$ is the number of times her neighbors select arm i. Using Lemma B.4, we will first find an upper bound for $n_i^1(t)$ for a suboptimal arm i.

Lemma B.4 states that at least one of the three inequalities (68a), (68b) and (68c) must be true. If Equation (68c) holds, then from (73) we obtain

$$n_i^1(t) \leq \frac{2\alpha}{\triangle_i^2} \ln t - n_i'(t). \tag{74}$$

Let U be the maximum of right hand side of (74) for $t = 1, .., T$:

$$U = \max \left\{ \max_{t=1,..,T} \left\{ \frac{2\alpha}{\triangle_i^2} \ln t - n_i'(t) \right\} \triangle_i, 0 \right\}, \tag{75}$$

as a result if the Equation (68c) holds for some instance k, then $n_i^1(k)$ is bounded by U, i.e.,

$$n_i^1(k) \leq U. \tag{76}$$

For bounding the regret we find an upper bound on the number of times we select a suboptimal arm i:

$$\mathbb{E}[n_i^1(T)] = \mathbb{E}\left[\sum_{t=1}^{T} 1_{a(t)=i}\right] = \mathbb{E}\left[\sum_{t=1}^{U} 1_{a(t)=i}\right] + \mathbb{E}\left[\sum_{t=U+1}^{T} 1_{a(t)=i}\right]. \tag{77}$$

Since $\mathbb{E}\left[\sum_{t=1}^{U} 1_{a(t)=i}\right] \leq U$, we can deduce

$$\mathbb{E}[n_i^1(T)] \leq U + \mathbb{E}\left[\sum_{t=U+1}^{T} 1_{a(t)=i}\right], \tag{78}$$

as we saw in (B.4), $1_{\{a(t)=i\}} = 1$ requires that at least one of the three equations (68a), (68b) and (68c) is true. Assume the last time that (68c) is true is at time ζ, hence

$$n_i^1(\zeta) \leq \frac{2\alpha}{\triangle_i^2} \ln \zeta - n_i'(\zeta), \tag{79}$$

and since ζ is the last time that (68c) holds, we can upper bound $\mathbb{E}[n_i^1(T)]$ by

$$\mathbb{E}[n_i^1(T)] \leq \frac{2\alpha}{\triangle_i^2} \ln \zeta - n_i'(\zeta) + \mathbb{E}\left[\sum_{t=\zeta+1}^{T} 1_{\{\text{(68a) } or \text{ (68b) } is\ true\ and \text{ (68c) } is\ false\}}\right]. \tag{80}$$

According to the definition of U in (75) and Equation (80) we have

$$\mathbb{E}[n_i^1(T)] \leq U + \mathbb{E}\left[\sum_{t=U+1}^{T} 1_{\{\text{(68a) } or \text{ (68b) } is\ true\ and \text{ (68c) } is\ false\}}\right] \tag{81}$$

$$\leq U + \sum_{t=U+1}^{T} \mathbb{P}[\text{(68a) is true}] + \mathbb{P}[\text{(68b) is true}]. \tag{82}$$

It suffices to bound the probability (68a) and (68b):

$$\mathbb{P}\left[\text{(68a) is true}\right] = \sum_{n_i(t)} \mathbb{P}\left[\text{(68a) is true}|n_i(t)\right] \cdot \mathbb{P}\left[n_i(t)\right]. \tag{83}$$

Recall that

$$\mathbb{P}\left[\text{(68a) is true}|n_i(t)\right] \leq \frac{1}{t^\alpha}. \tag{84}$$

Plugging (84) in (83) yields that

$$\mathbb{P}\left[\text{(68a) is true}\right] \leq \frac{1}{t^\alpha} \sum_{n_i(t)} \mathbb{P}\left[n_i(t)\right] = \frac{1}{t^\alpha}. \tag{85}$$

Then we take integral of $\frac{1}{t^\alpha}$ for t from 1 to T, which is smaller than $\frac{\alpha}{2(\alpha-2)}$. The same upper bound holds for (68b).

Thus, the regret is

$$R \leq \sum_{i,\triangle_i>0} \left(\max\left\{\max_{t=1,..,T}\left\{\frac{2\alpha}{\triangle_i^2}\ln t - n_i'(t)\right\}\triangle_i, 0\right\} + \frac{\alpha}{\alpha-2}\right) \tag{86}$$

as desired. □

B.1 UCBN on Complete Graphs

In this section, we analyze the regret in a complete graph when all agents use UCBN. The following curious lemma will assist in the proof.

Lemma B.5. *Given a complete graph of agents, if all agents use* UCBN *with a deterministic common tie breaking scheme, then in every time step all agents select the same action.*

Proof. Since the graph is complete, all agents see the rewards of other agents at every time step; hence the sample means $\widehat{s}_i(t)$ and number of samples $n_i(t)$ at time t are the same for all agents. Furthermore, every agent selects an arm according to criteria (1) and a common deterministic tie breaking rule. Therefore, the arm selected at time t will be same for all agents. □

Proof of Corollary B.2. Let

$$U = \left\lceil \frac{2\alpha \ln T}{b \cdot \triangle_i^2} \right\rceil. \tag{87}$$

We bound the number of times action i other than the best arm is selected. Following the proof of Theorem 3.1,

$$\mathbb{E}[n_i^1(T)] \leq U + \sum_{t=U+1}^{T} \mathbb{P}[\text{(68a) is true}] + \mathbb{P}[\text{(68b) is true}]. \tag{88}$$

The upper bound of the probabilities (68a) and (68b) are same as before. Hence, the regret bound is

$$R \leq \sum_{i,\triangle_i>0} \left(\frac{2\alpha}{b \cdot \triangle_i} \ln T + \frac{\alpha}{\alpha - 2} \right). \tag{89}$$

□

B.2 Lower Bound

The lower bound for UCB [LR85] is

$$\lim_{T \longrightarrow +\infty} \inf \frac{R}{\ln T} \geq \sum_{i,\triangle_i>0} \frac{\triangle_i}{kl(\mu_i, \mu^\star)}. \tag{90}$$

Our proof again follows the same template.

Proof of Theorem B.3. As in [LR85], we assume the rewards are drawn from a Bernoulli distribution. From their proof it follows that the expected number of times that a suboptimal arm must be selected in order to distinguish between best arm and other arms is at least

$$\mathbb{E}[n_i(T)] + c_T(i) \geq (1 + o(1)) \frac{1-\varepsilon}{1+\varepsilon} \frac{\ln T}{kl(\mu_i, \mu^\star)}. \tag{91}$$

where the second term is the information coming from the neighbors (that is, the number of times neighbors selected arm i up to time t), $\mu^\star$ is the mean of the best arm, and $kl(p, q)$ is the Kullback-Leibler divergence between a Bernoulli variable with parameter p and a Bernoulli variable with parameter q, defined to be

$$kl(p, q) \stackrel{\text{def}}{=} p \ln \left(\frac{p}{q} \right) + (1-p) \ln \left(\frac{1-p}{1-q} \right). \tag{92}$$

As the number of rounds increases, ε can be taken to be smaller. As T goes to infinity, ε can be taken zero; as a result we can write the following lower bound for the regret

$$\lim_{T \longrightarrow +\infty} \inf \frac{R + \sum_{i,\Delta_i>0} c_T(i) \Delta_i}{\ln T} \geq \sum_{i,\triangle_i>0} \frac{\triangle_i}{kl(\mu_i, \mu^\star)}. \tag{93}$$

Applying Pinkser's inequality yields the lower bound. □